Weak Subgame Perfect Equilibria and their Application to Quantitative Reachability*

Thomas Brihaye[1], Véronique Bruyère[2], Noémie Meunier[†2], and Jean-François Raskin[‡3]

1 Département de mathématique, Université de Mons (UMONS) Mons, Belgium
2 Département d'informatique, Université de Mons (UMONS) Mons, Belgium
3 Département d'informatique, Université Libre de Bruxelles (U.L.B.) Brussels, Belgium

Abstract

We study n-player turn-based games played on a finite directed graph. For each play, the players have to pay a cost that they want to minimize. Instead of the well-known notion of Nash equilibrium (NE), we focus on the notion of subgame perfect equilibrium (SPE), a refinement of NE well-suited in the framework of games played on graphs. We also study natural variants of SPE, named weak (resp. very weak) SPE, where players who deviate cannot use the full class of strategies but only a subclass with a finite number of (resp. a unique) deviation step(s).

Our results are threefold. Firstly, we characterize in the form of a Folk theorem the set of all plays that are the outcome of a weak SPE. We also establish a weaker version of this theorem for SPEs. Secondly, for the class of quantitative reachability games, we prove the existence of a finite-memory SPE and provide an algorithm for computing it (only existence was known with no information regarding the memory). Moreover, we show that the existence of a constrained SPE, i.e. an SPE such that each player pays a cost less than a given constant, can be decided. The proofs rely on our Folk theorem for weak SPEs (which coincide with SPEs in the case of quantitative reachability games) and on the decidability of MSO logic on infinite words. Finally with similar techniques, we provide a second general class of games for which the existence of a (constrained) weak SPE is decidable.

1998 ACM Subject Classification B.6.3 [design aids]: automatic synthesis; F.1.2 [Modes of computation]: interactive and reactive computation

Keywords and phrases multi-player games on graphs, quantitative objectives, Nash equilibrium, subgame perfect equilibrium, quantitative reachability

1 Introduction

Two-player zero-sum infinite duration games played on graphs are a mathematical model used to formalize several important problems in computer science. Reactive system synthesis is one such important problem. In this context, see e.g. [13], the vertices and the edges of the graph represent the states and the transitions of the system; one player models the system to synthesize, and the other player models the (uncontrollable) environment of the system.

* This work has been partly supported by European project Cassting (FP7-ICT-601148).
† Author supported by F.R.S.-FNRS fellowship.
‡ Author supported by ERC Starting Grant (279499: inVEST).

Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In the classical setting, the objectives of the two players are opposite, i.e. the environment is *adversarial*. Modeling the environment as fully adversarial is usually a bold abstraction of reality and there are recent works that consider the more general setting of non zero-sum games which allow to take into account the different objectives of each player. In this latter setting the environment has its own objective which is most often *not* the negation of the objective of the system. The concept of *Nash equilibrium* (NE) [11] is central to the study of non zero-sum games and can be applied to the general setting of n player games. A strategy profile is a NE if no player has an incentive to deviate unilaterally from his strategy, since he cannot strictly improve on the outcome of the strategy profile by changing his strategy only.

However in the context of sequential games (such as games played on graphs), it is well-known that NEs present a serious weakness: a NE allows for *non-credible threats* that rational players should not carry out [15]. As a consequence, for sequential games, the notion of NE has been strengthened into the notion of *subgame perfect equilibrium* (SPE): a strategy profile is an SPE if it is a NE in all the subgames of the original game. While the notion of SPE is rather well understood for finite state game graphs with ω-regular objectives or for games in finite extensive form (finite game trees), less is known for game graphs with *quantitative objectives* in which players encounter costs that they want to minimize, like in classical quantitative objectives such as mean-payoff, discounted sum, or quantitative reachability.

Several natural and important questions arise for such games: Can we decide the existence of an SPE, and more generally the *constrained* existence of an SPE (i.e. an SPE in which each player encounters a cost less than some fixed value)? Can we compute such SPEs that use finite-memory strategies only? Whereas several work has studied what are the hypotheses to impose on games in a way to guarantee the existence of an SPE, the previous algorithmic questions are still widely open. In this article, we provide progresses in the understanding of the notion of SPE. We first establish Folk theorems that characterize the possible outcomes of SPEs in quantitative games for SPEs and their variants. We then derive from this characterization interesting algorithms and information on the strategies for two important classes of quantitative games. Our contributions are detailed in the next paragraph.

Contributions First, we formalize a notion of *deviation step* from a strategy profile that allows us to define two natural variants of NEs. While a NE must be resistant to the unilateral deviation of one player for any number of deviation steps, a *weak* (resp. *very weak*) NE must be resistant to the unilateral deviation of one player for any *finite* number of (resp. a unique) deviation step(s). Then we use those variants to define the corresponding notions of *weak* and *very weak* SPE. The latter notion is very close to the one-step deviation property [12]. Any very weak SPE is also a weak SPE, and there are games for which there exists a weak SPE but no SPE. Also, for games with upper-semicontinuous cost functions and for games played on finite game trees, the three notions are equivalent.

Second, we characterize in the form of a Folk theorem all the possible outcomes of weak SPEs. The characterization is obtained starting from all possible plays of the game and the application of a nonincreasing operator that removes plays that cannot be outcome of a weak SPE. We show that the limit of the nonincreasing chain of sets always exists and contains exactly all the possible outcomes of weak SPEs. Furthermore, we show how for each such outcome, we can associate a strategy profile that generates it and which is a weak SPE. Using a variant of the techniques developed for weak SPEs, we also obtain a weaker version of this Folk theorem for SPEs.

Additionally, to illustrate the potential of our Folk theorem, we show how it can be

refined and used to answer open questions about two classes of quantitative games. The first class of games that we consider are *quantitative reachability games*, such that each player aims at reaching his own set of target states as soon as possible. As the cost functions in those games are continuous, our Folk theorem characterizes precisely the outcomes of SPEs and not only weak SPEs. In [1, 6], it has been shown that quantitative reachability games always have SPEs. The proof provided for this theorem is non constructive since it relies on topological arguments. Here, we strengthen this existential result by proving that there always exists, not only an SPE but, a *finite-memory* SPE. Furthermore, we provide an algorithm to construct such a finite memory SPE. This algorithm is based on a constructive version of our Folk Theorem for the class of quantitative reachability games: we show that the nonincreasing chain of sets of potential outcomes stabilizes after a finite number of steps and that each intermediate set is an ω-regular set that can be effectively described using MSO sentences. The second class of games that we consider is the class of games with cost functions that are *prefix-independent*, whose range of values is *finite*, and for which each value has an ω*-regular* pre-image. For this general class of games, with similar techniques as for quantitative reachability games, we show how to construct an effective representation of all possible outcomes compatible with a weak SPE, and consequently that the existence of a weak SPE is decidable. In those two applications, we show that our construction also allow us to answer the question of existence of a constrained (weak) SPE, i.e. a (weak) SPE in which players pays a cost which is bounded by a given value.

Related work The concept of SPE has been first introduced and studied by the game theory community. The notion of SPE has been first introduced by Kuhn in finite extensive form games [9]. For such games, backward induction can be used to prove that there always exist an SPE. By inspecting the backward induction proof, it is not difficult to realize that the notion of very weak SPE and SPE are equivalent in this context.

SPEs for infinite trees defined as the unfolding of finite graphs with *qualitative*, i.e. win-lose, ω-regular objectives, have been studied by Ummels in [18]: it is proved that such games always have an SPE, and that the existence of a constrained SPE is decidable.

In [8], Klimos et al. provide an effective representation of the outcomes of NEs in concurrent priced games by constructing a Büchi automaton accepting precisely the language of outcomes of all NEs satisfying a bound vector. The existence of NEs in quantitative games played on graphs is studied in [2]; it is shown that for a large class of games, there always exists a finite-memory NE. This result is extended in [3] for two-player games and secure equilibria (a refinement of NEs); additionally the constrained existence problem for secure equilibria is also shown decidable for a large range of cost functions. None of these references consider SPEs.

In [5], the authors prove that for quantitative games with cost functions that are upper-semicontinus and with finite range, there always exits an SPE. This result also relies on a nonincreasing chain of sets of possible outcomes of SPEs. The main differences with our work is that we obtain a Folk theorem that *characterize* all possible outcomes of weak SPEs with no restriction on the cost functions. Moreover we have shown that our Folk theorem can be made effective for two classes of quantitative games of interest. Effectiveness issues are not considered in [5]. Prior to this work, Mertens shows in [10] that if the cost functions are bounded and Borel measurable then there always exists an ϵ-NE. In [6], Fudenberg et al. show that if the cost functions are all continous, then there always exists an SPE. Those works were recently extended in [14] by Le Roux and Pauly.

Organization of the article In Section 2, we present the notions of quantitative game, classical NE and SPE, and their variants. In Section 3, we propose and prove our Folk

Theorems for weak SPEs and for SPEs. In Section 4, we provide an algorithm for computing a finite-memory SPE for quantitative reachability games, and a second algorithm to decide the constrained existence of an SPE for this class of games. In Section 5, we show that the existence of a (constrained) weak SPE is decidable for another class of games. A conclusion and future work are given in the last section.

2 Preliminaries and Variants of Equilibria

In this section, we recall the notions of quantitative game, Nash equilibrium, and subgame perfect equilibrium. We also introduce variants of Nash and subgame perfect equilibria, and compare them with the classical notions.

2.1 Quantitative Games

We consider multi-player turn-based non zero-sum quantitative games in which, for each infinite play, players pay a cost that they want to minimize.[1]

▶ **Definition 1.** A *quantitative game* is a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \bar{\lambda})$ where:

- Π is a finite set of players,
- V is a finite set of vertices,
- $(V_i)_{i\in\Pi}$ is a partition of V such that V_i is the set of vertices controlled by player $i \in \Pi$,
- $E \subseteq V \times V$ is a set of edges, such that[2] for all $v \in V$, there exists $v' \in V$ with $(v, v') \in E$,
- $\bar{\lambda} = (\lambda_i)_{i\in\Pi}$ is a cost function such that $\lambda_i : V^\omega \to \mathbb{R} \cup \{+\infty\}$ is player i cost function.

A *play* of $\mathcal{G}$ is an infinite sequence $\rho = \rho_0\rho_1 \ldots \in V^\omega$ such that $(\rho_i, \rho_{i+1}) \in E$ for all $i \in \mathbb{N}$. *Histories* of $\mathcal{G}$ are finite sequences $h = h_0 \ldots h_n \in V^+$ defined in the same way. The *length* $|h|$ of h is the number n of its edges. We denote by $\mathrm{First}(h)$ (resp. $\mathrm{Last}(h)$) the first vertex h_0 (resp. last vertex h_n) of h. Usually histories are non-empty, but in specific situations it will be useful to consider the empty history ϵ. The set of all histories (ended by a vertex in V_i) is denoted by Hist (by Hist_i). A *prefix* (resp. *suffix*) of a play ρ is a finite sequence $\rho_0 \ldots \rho_n$ (resp. infinite sequence $\rho_n\rho_{n+1} \ldots$) denoted by $\rho_{\le n}$ or $\rho_{<n+1}$ (resp. $\rho_{\ge n}$). We use notation $h < \rho$ when a history h is prefix of a play ρ. Given two distinct plays ρ and ρ', their longest common prefix is denoted by $\rho\,\widehat{}\,\rho'$.

When an initial vertex $v_0 \in V$ is fixed, we call $(\mathcal{G}, v_0)$ an *initialized* quantitative game. A play (resp. a history) of $(\mathcal{G}, v_0)$ is a play (resp. history) of $\mathcal{G}$ starting in v_0. The set of histories $h \in \mathrm{Hist}$ (resp. $h \in \mathrm{Hist}_i$) with $\mathrm{First}(h) = v_0$ is denoted by $\mathrm{Hist}(v_0)$ (resp. $\mathrm{Hist}_i(v_0)$). In the figures of this article, we will often *unravel* the graph of the game $(\mathcal{G}, v_0)$ from the initial vertex v_0, which ends up in an infinite tree.

Given a play $\rho \in V^\omega$, its *cost* is given by $\bar{\lambda}(\rho) = (\lambda_i(\rho))_{i\in\Pi}$. In this article, we are particularly interested in quantitative reachability games in which $\lambda_i(\rho)$ is equal to the number of edges to reach a given set of vertices.

▶ **Definition 2.** A *quantitative reachability game* is a quantitative game $\mathcal{G}$ such that the cost function $\bar{\lambda} : V^\omega \to (\mathbb{N} \cup \{+\infty\})^\Pi$ is defined as follows. Each player i has a *target set* $T_i \subseteq V$, and for each play $\rho = \rho_0\rho_1 \ldots$ of $\mathcal{G}$, the cost $\lambda_i(\rho)$ is the least index n such that $\rho_n \in T_i$ if it exists, and $+\infty$ otherwise.

[1] Alternatively, players could receive a payoff that they want to maximize.
[2] Each vertex has at least one outgoing edge.

Notice that the cost function $\bar{\lambda}$ of a quantitative game is often defined from $|\Pi|$-uples of weights labeling the edges of the game. For instance, in inf games, $\lambda_i(\rho)$ is equal to the infimum of player i weights seen along ρ. Some other classical examples are liminf, limsup, mean-payoff, and discounted sum games [4]. In case of quantitative reachability on graphs with weighted edges, the cost $\lambda_i(\rho)$ for player i is replaced by the sum of the weights seen along ρ until his target set is reached. We do not consider this extension here. Notice that when weights are positive integers, replacing each edge with cost c by a path of length c composed of c new edges allows to recover Definition 2.

Let us recall the notions of prefix-independent, continuous, and lower- (resp. upper-) semicontinuous cost functions. Since V is endowed with the discrete topology, and thus V^ω with the product topology, a sequence of plays $(\rho_n)_{n\in\mathbb{N}}$ converges to a play $\rho = \lim_{n\to\infty} \rho_n$ if every prefix of ρ is prefix of all ρ_n except, possibly, of finitely many of them.

▶ **Definition 3.** Let λ_i be a player i cost function. Then

- λ_i is *prefix-independent* if $\lambda_i(h\rho) = \lambda_i(\rho)$ for any history h and play ρ.
- λ_i is *continuous* if whenever $\lim_{n\to\infty} \rho_n = \rho$, then $\lim_{n\to\infty} \lambda_i(\rho_n) = \lambda_i(\rho)$.
- λ_i *upper-semicontinuous* (resp. *lower-semicontinuous*) if whenever $\lim_{n\to\infty} \rho_n = \rho$, then $\limsup_{n\to\infty} \lambda_i(\rho_n) \leq \lambda_i(\rho)$ (resp. $\liminf_{n\to\infty} \lambda_i(\rho_n) \geq \lambda_i(\rho)$).

For instance, the cost functions used in liminf and mean-payoff games are prefix-independent, contrarily to the case of inf games. Clearly, if λ_i is continuous, then it is upper- and lower-semicontinuous. For instance, the cost functions of liminf and mean-payoff games are neither upper-semicontinuous nor lower-semicontinuous, whereas cost functions of discounted sum games are continuous. The cost functions λ_i, $i \in \Pi$, used in quantitative reachability games can be transformed into continuous ones as follows [1]: $\lambda'_i(\rho) = 1 - \frac{1}{\lambda_i(\rho)+1}$ if $\lambda_i(\rho) < +\infty$, and $\lambda'_i(\rho) = 1$ otherwise.

2.2 Strategies and Deviations

A *strategy* σ for player $i \in \Pi$ is a function $\sigma : \mathrm{Hist}_i \to V$ assigning to each history[3] $hv \in \mathrm{Hist}_i$ a vertex $v' = \sigma(hv)$ such that $(v, v') \in E$. In an initialized game $(\mathcal{G}, v_0)$, σ is restricted to histories starting with v_0. A player i strategy σ is *positional* if it only depends on the last vertex of the history, i.e. $\sigma(hv) = \sigma(v)$ for all $hv \in \mathrm{Hist}_i$. It is a *finite-memory* strategy if it needs only finite memory of the history (recorded by a finite strategy automaton, also called a Moore machine). A play ρ is *consistent* with a player i strategy σ if $\rho_{k+1} = \sigma(\rho_{\leq k})$ for all k such that $\rho_k \in V_i$. A *strategy profile* of $\mathcal{G}$ is a tuple $\bar{\sigma} = (\sigma_i)_{i\in\Pi}$ of strategies, where each σ_i is a player i strategy. It is called *positional* (resp. *finite-memory*) if all σ_i, $i \in \Pi$, are positional (resp. finite-memory). Given an initial vertex v_0, such a strategy profile determines a unique play of $(\mathcal{G}, v_0)$ that is consistent with all the strategies. This play is called the *outcome* of $\bar{\sigma}$ and is denoted by $\langle\bar{\sigma}\rangle_{v_0}$.

Given σ_i a player i strategy, we say that player i *deviates* from σ_i if he does not stick to σ_i and prefers to use another strategy σ'_i. Let $\bar{\sigma}$ be a strategy profile. When all players stick to their strategy σ_i except player i that shifts to σ'_i, we denote by (σ'_i, σ_{-i}) the derived strategy profile, and by $\langle\sigma'_i, \sigma_{-i}\rangle_{v_0}$ its outcome in $(\mathcal{G}, v_0)$. In the next definition, we introduce the notion of deviation step of a strategy σ'_i from a given strategy profile $\bar{\sigma}$.

[3] In this article we often write a history in the form hv with $v \in V$ to emphasize that v is the last vertex of this history.

▶ **Definition 4.** Let $(\mathcal{G}, v_0)$ be an initialized game, $\bar{\sigma}$ be a strategy profile, and σ'_i be a player i strategy. We say that σ'_i has a *hv-deviation step* from $\bar{\sigma}$ for some history $hv \in \mathrm{Hist}_i(v_0)$ with $v \in V_i$, if

$$hv < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0} \text{ and } \sigma_i(hv) \neq \sigma'_i(hv).$$

Notice that the previous definition requires that hv is a prefix of the outcome $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$; it says nothing about σ'_i outside of this outcome. A strategy σ'_i can have a finite or an infinite number of deviation steps in the sense of Definition 4. A strategy with three deviation steps is depicted in Figure 1 (left) such that each $h_k v_k$-deviation step from $\bar{\sigma}$, $1 \leq k \leq 3$, is highlighted with a dashed edge. We will come back to this figure later on.

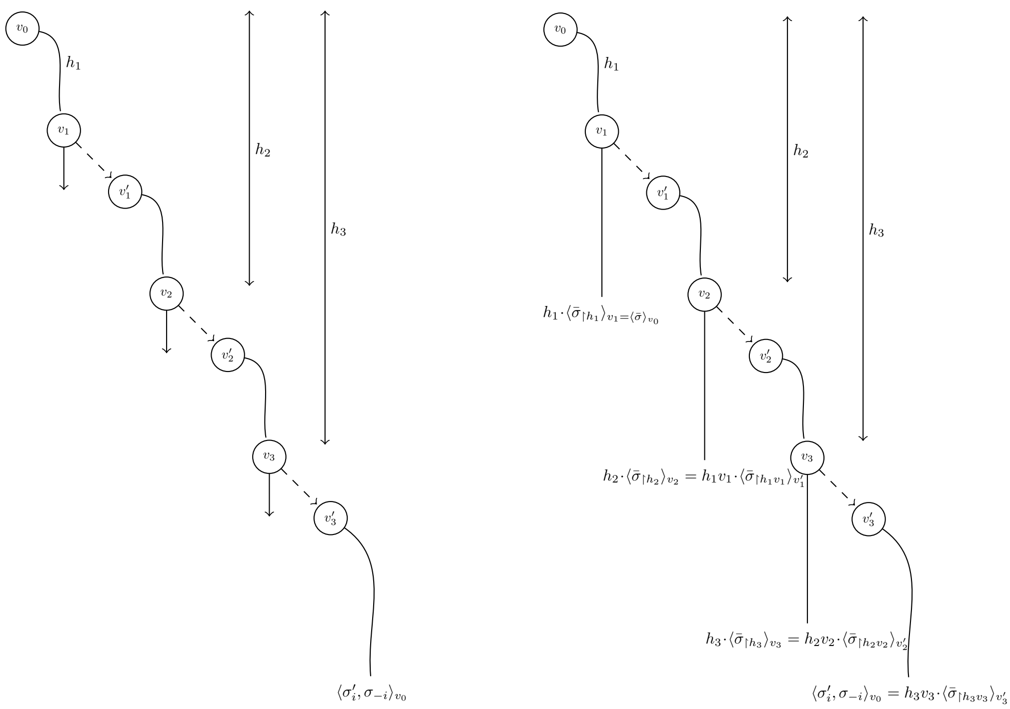

Figure 1 A strategy σ'_i with a finite number of deviation steps.

In light of Definition 4, we introduce the following classes of strategies.

▶ **Definition 5.** Let $(\mathcal{G}, v_0)$ be an initialized game, and $\bar{\sigma}$ be a strategy profile.

- A strategy σ'_i is *finitely deviating* from $\bar{\sigma}$ if it has a finite number of deviation steps from $\bar{\sigma}$.
- It is *one-shot deviating* from $\bar{\sigma}$ if it has a v_0-deviation step from $\bar{\sigma}$, and no other deviation step.

In other words, a strategy σ'_i is finitely deviating from $\bar{\sigma}$ if there exists a history $hv < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$ such that for all $h'v'$, $hv \leq h'v' < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$, we have $\sigma'_i(h'v') = \sigma_i(h'v')$ (σ'_i acts as σ_i from hv along $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$). The strategy σ'_i is one-shot deviating from $\bar{\sigma}$ if it differs from σ_i at the initial vertex v_0, and after v_0 acts as σ_i along $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$. As for Definition 4, the previous definition says nothing about σ'_i outside of $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$. Clearly any one-shot deviating strategy is finitely deviating. The strategy of Figure 1 is finitely deviating but not one-shot deviating.

2.3 Nash and Subgame Perfect Equilibria, and Variants

In this paper, we focus on subgame perfect equilibria and their variants. Let us first recall the classical notion of Nash equilibrium. A strategy profile $\bar{\sigma}$ in an initialized game is a Nash equilibrium if no player has an incentive to deviate unilaterally from his strategy, since he cannot strictly decrease his cost when using any other strategy.

▶ **Definition 6.** Given an initialized game $(\mathcal{G}, v_0)$, a strategy profile $\bar{\sigma} = (\sigma_i)_{i\in\Pi}$ of $(\mathcal{G}, v_0)$ is a *Nash equilibrium (NE)* if for all players $i \in \Pi$, for all player i strategies σ'_i, we have $\lambda_i(\langle\sigma'_i, \sigma_{-i}\rangle_{v_0}) \geq \lambda_i(\langle\bar{\sigma}\rangle_{v_0})$.

We say that a player i strategy σ'_i is a *profitable deviation* for i w.r.t. $\bar{\sigma}$ if $\lambda_i(\langle\sigma'_i, \sigma_{-i}\rangle_{v_0}) < \lambda_i(\langle\bar{\sigma}\rangle_{v_0})$. Therefore $\bar{\sigma}$ is a NE if no player has a profitable deviation w.r.t. $\bar{\sigma}$.

Let us propose the next variants of NE.

▶ **Definition 7.** Let $(\mathcal{G}, v_0)$ be an initialized game. A strategy profile $\bar{\sigma}$ is a *weak NE* (resp. *very weak NE*) in $(\mathcal{G}, v_0)$ if, for each player $i \in \Pi$, for each finitely deviating (resp. one-shot deviating) strategy σ'_i of player i, we have $\lambda_i(\langle\sigma'_i, \sigma_{-i}\rangle_{v_0}) \geq \lambda_i(\langle\bar{\sigma}\rangle_{v_0})$.

▶ **Example 8.** Consider the two-player quantitative game depicted in Figure 2. Circle (resp. square) vertices are player 1 (resp. player 2) vertices. The edges are labeled by couples of weights such that weights $(0, 0)$ are not specified. For each player i, the cost $\lambda_i(\rho)$ of a play ρ is the weight of its ending loop. In this simple game, each player i have two positional strategies that are respectively denoted by σ_i and σ'_i (see Figure 2).

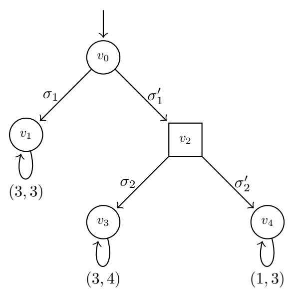

Figure 2 A simple two-player quantitative game

The strategy profile (σ_1, σ'_2) is not a NE since σ'_1 is a profitable deviation for player 1 w.r.t. (σ_1, σ'_2) (player 1 pays cost 1 instead of cost 3). This strategy profile is neither a weak NE nor a very weak NE because in this simple game, player 1 can only deviate from σ_1 by using the one-shot deviating strategy σ'_1. On the contrary, the strategy profile (σ_1, σ_2) is a NE with outcome $v_0 v_1^\omega$ of cost $(3, 3)$. It is also a weak NE and a very weak NE.

By definition, any NE is a weak NE, and any weak NE is a very weak NE. The contrary is false: in the previous example, (σ'_1, σ_2) is a very weak NE, but not a weak NE. We will see later an example of game with a weak NE that is not an NE (see Example 12).

The notion of subgame perfect equilibrium is a refinment of NE. In order to define it, we need to introduce the following notions. Given a quantitative game $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \bar{\lambda})$ and a history h of $\mathcal{G}$, we denote by $\mathcal{G}_{\restriction h}$ the game $\mathcal{G}_{\restriction_h} = (\Pi, V, (V_i)_{i\in\Pi}, E, \bar{\lambda}_{\restriction_h})$ where

$$\bar{\lambda}_{\restriction h}(\rho) = \bar{\lambda}(h\rho)$$

for any play of $\mathcal{G}_{\restriction h}$[4], and we say that $\mathcal{G}_{\restriction h}$ is a *subgame* of $\mathcal{G}$. Given an initialized game $(\mathcal{G}, v_0)$, and a history $hv \in \mathrm{Hist}(v_0)$, the initialized game $(\mathcal{G}_{\restriction h}, v)$ is called the subgame of $(\mathcal{G}, v_0)$ with history hv. Notice that $(\mathcal{G}, v_0)$ can be seen as a subgame of itself with history hv_0 such that $h = \epsilon$. Given a player i strategy σ in $(\mathcal{G}, v_0)$, we define the strategy $\sigma_{\restriction h}$ in $(\mathcal{G}_{\restriction h}, v)$ as $\sigma_{\restriction h}(h') = \sigma(hh')$ for all histories $h' \in \mathrm{Hist}_i(v)$. Given a strategy profile $\bar{\sigma} = (\sigma_i)_{i\in\Pi}$, we use notation $\bar{\sigma}_{\restriction h}$ for $(\sigma_{i\restriction h})_{i\in\Pi}$, and $\langle\bar{\sigma}_{\restriction h}\rangle_v$ is its outcome in the subgame $(\mathcal{G}_{\restriction h}, v)$.

We can now recall the classical notion of subgame perfect equilibrium: it is a strategy profile in an initialized game that induces a NE in each of its subgames. In particular, a subgame perfect equilibrium is a NE.

▶ **Definition 9.** Given an initialized game $(\mathcal{G}, v_0)$, a strategy profile $\bar{\sigma}$ of $(\mathcal{G}, v_0)$ is a *subgame perfect equilibrium (SPE)* if $\bar{\sigma}_{\restriction h}$ is a NE in $(\mathcal{G}_{\restriction h}, v)$, for every history $hv \in \mathrm{Hist}(v_0)$.

As for NE, we propose the next variants of SPE.

▶ **Definition 10.** Let $(\mathcal{G}, v_0)$ be an initialized game. A strategy profile $\bar{\sigma}$ is a *weak SPE* (resp. *very weak SPE*) if $\bar{\sigma}_{\restriction h}$ is a weak NE (resp. very weak NE) in $(\mathcal{G}_{\restriction h}, v)$, for all histories $hv \in \mathrm{Hist}(v_0)$.

▶ **Example 11.** We come back to the game depicted in Figure 2. We have seen before that the strategy profile (σ_1, σ_2) is a NE. However it is not an SPE. Indeed consider the subgame $(\mathcal{G}_{\restriction v_0}, v_2)$ of $(\mathcal{G}, v_0)$ with history v_0v_2. In this subgame, σ'_2 is a profitable deviation for player 2. One can easily verify that the strategy profile (σ'_1, σ'_2) is an SPE, as well as a weak SPE and a very weak SPE, due to the simple form of the game.

The previous example is too simple to show the differences between classical SPEs and their variants. The next example presents a game with a (very) weak SPE but no SPE.

▶ **Example 12.** Consider the initialized two-player game $(\mathcal{G}, v_0)$ in Figure 3. The edges are labeled by couples of weights, and for each player i the cost $\lambda_i(\rho)$ of a play ρ is the unique weight seen in its ending cycle. With this definition, $\lambda_i(\rho)$ can also be seen as either the mean-payoff, or the liminf, or the limsup, of the weights of ρ. It is known that this game has no SPE [16].

Let us show that the positional strategy profile $\bar{\sigma}$ depicted with thick edges is a very weak SPE. Due to the simple form of the game, only two cases are to be treated. Consider the subgame $(\mathcal{G}_{\restriction h}, v_0)$ with $h \in (v_0v_1)^*$, and the one-shot deviating strategy σ'_1 of player 1 such that $\sigma'_1(v_0) = v_2$. Then $\langle\bar{\sigma}_{\restriction h}\rangle_{v_0} = v_0v_1v_3^\omega$ and $\langle\sigma'_1, \sigma_{2\restriction h}\rangle_{v_0} = v_0v_2^\omega$, showing that σ'_1 is not a profitable deviation for player 1. One also checks that in the subgame $(\mathcal{G}_{\restriction h}, v_1)$ with $h \in (v_0v_1)^*v_0$, the one-shot deviating strategy σ'_2 of player 2 such that $\sigma'_2(v_1) = v_0$ is not profitable for him.

Similarly, one can prove that $\bar{\sigma}$ is a weak SPE (see also Proposition 13 hereafter). Notice that $\bar{\sigma}$ is not an SPE. Indeed the strategy σ'_2 such that $\sigma'_2(hv_1) = v_0$ for all h, is a profitable deviation for player 2 in $(\mathcal{G}, v_0)$. This strategy is (of course) not finitely deviating. Finally notice that $\bar{\sigma}$ is a weak NE that is not an NE.

From Definition 10, any SPE is a weak SPE, and any weak SPE is a very weak SPE. The next proposition states that weak SPE and very weak SPE are equivalent notions, but this is no longer true for SPE and weak SPE as shown previously by Example 12.

[4] In this article, we will always use notation $\bar{\lambda}(h\rho)$ instead of $\bar{\lambda}_{\restriction h}(\rho)$.

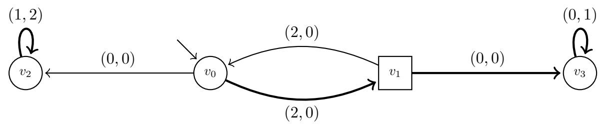

Figure 3 A two-player game with a (very) weak SPE and no SPE. For each player, the cost of a play is his unique weight seen in the ending cycle.

▶ **Proposition 13.**
- *Let $(\mathcal{G}, v_0)$ be an initialized game, and $\bar{\sigma}$ be a strategy profile. Then $\bar{\sigma}$ is a weak SPE iff $\bar{\sigma}$ is a very weak SPE.*
- *There exists an initialized game $(\mathcal{G}, v_0)$ with a weak SPE but no SPE.*

Before proving this proposition, we would like to come back to the definition of deviation step (Definition 4), and explain it now with the concept of subgame. Given an initialized game $(\mathcal{G}, v_0)$, a strategy profile $\bar{\sigma}$, and a player i strategy σ'_i, we recall that σ'_i has a hv-deviation step from $\bar{\sigma}$ for some $hv \in \mathrm{Hist}_i(v_0)$, if $hv < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$ and $\sigma_i(hv) \neq \sigma'_i(hv)$. Equivalently, σ'_i has a hv-deviation step from $\bar{\sigma}$ iff

$$h \cdot \langle \bar{\sigma}_{\restriction h} \rangle_v \,\hat{}\, \langle \sigma'_i, \sigma_{-i} \rangle_{v_0} = hv.$$

This alternative vision of deviation step is depicted in Figure 1 (right) for a strategy with three deviation steps. For instance, for history $h_2 v_2$, we have $h_2 v_2 < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$ and $\sigma_i(h_2 v_2) \neq \sigma'_i(h_2 v_2)$, or equivalently $h_2 \cdot \langle \bar{\sigma}_{\restriction h_2} \rangle_{v_2} \,\hat{}\, \langle \sigma'_i, \sigma_{-i} \rangle_{v_0} = h_2 v_2$. Notice that there is no intermediate deviation step between the $h_1 v_1$-deviation step and the $h_2 v_2$-deviation step since $h_2 \cdot \langle \bar{\sigma}_{\restriction h_2} \rangle_{v_2} = h_1 v_1 \cdot \langle \bar{\sigma}_{\restriction h_1 v_1} \rangle_{v'_1}$ as indicated in the figure. Similarly the $h_1 v_1$-deviation step is the first one because $h_1 \cdot \langle \bar{\sigma}_{\restriction h_1} \rangle_{v_1} = \langle \bar{\sigma} \rangle_{v_0}$, and the $h_3 v_3$-deviation step is the third one because $h_3 \cdot \langle \bar{\sigma}_{\restriction h_3} \rangle_{v_3} = h_2 v_2 \cdot \langle \bar{\sigma}_{\restriction h_2 v_2} \rangle_{v'_2}$. The latter deviation step is the last one because $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0} = h_3 v_3 \cdot \langle \bar{\sigma}_{\restriction h_3 v_3} \rangle_{v'_3}$.

Proof of Proposition 13. This proof is based on arguments from the one-step deviation property used to prove Kuhn's theorem [9]. Let $\bar{\sigma}$ be a very weak SPE, and let us prove that it is a weak SPE. As a contradiction, assume that there exists a subgame $(\mathcal{G}_{\restriction h}, v)$ such that the strategy profile $\bar{\sigma}_{\restriction h}$ is not a weak NE. This means that there exists a player i strategy σ'_i in $(\mathcal{G}_{\restriction h}, v)$ such that σ'_i is finitely deviating from $\bar{\sigma}_{\restriction h}$ and

$$\lambda_i(h\rho) > \lambda_i(h\rho'), \tag{1}$$

where $\rho = \langle \bar{\sigma}_{\restriction h} \rangle_v$ and $\rho' = \langle \sigma'_i, \sigma_{-i \restriction h} \rangle_v$. Let us consider such a strategy σ'_i with a minimum number n of deviation steps from $\langle \bar{\sigma}_{\restriction h} \rangle_v$, and let $g_k v_k$, $1 \leq k \leq n$, be the histories in $\mathrm{Hist}_i(v)$ such that σ'_i has a $g_k v_k$-deviation step from $\langle \bar{\sigma}_{\restriction h} \rangle_v$. Let us consider the subgame $(\mathcal{G}_{\restriction h g_n}, v_n)$ (see Figure 4). In this subgame, $\sigma'_{i \restriction g_n}$ is not a profitable one-shot deviating strategy as $\bar{\sigma}$ is a very weak SPE. In other words, for $\varrho = \langle \bar{\sigma}_{\restriction h g_n} \rangle_{v_n}$ and $\varrho' = \langle \sigma'_{i \restriction g_n}, \sigma_{-i \restriction h g_n} \rangle_{v_n}$, we have

$$\lambda_i(h g_n \varrho) \leq \lambda_i(h g_n \varrho'). \tag{2}$$

Notice that $n \geq 2$. Indeed, if $n = 1$, then $\rho = g_1 \varrho$, $\rho' = g_1 \varrho'$, and $\lambda_i(h\rho) \leq \lambda_i(h\rho')$ by (2). Therefore σ'_i is not a profitable deviation in $(\mathcal{G}_{\restriction h}, v)$, in contradiction with its definition (1). We can thus construct a strategy τ'_i from σ'_i such that these two strategies are the same except in the subgame $(\mathcal{G}_{\restriction h g_n}, v_n)$ where $\tau'_{i \restriction g_n}$ and $\sigma_{i \restriction h g_n}$ coincide. In other words τ'_i has $n-1$ deviation steps from $\langle \bar{\sigma}_{\restriction h} \rangle_v$, that are exactly the $g_k v_k$-deviation steps, $1 \leq k \leq n-1$, of σ'_i. Moreover, in the subgame $(\mathcal{G}_{\restriction h}, v)$, we have $\langle \tau'_i, \sigma_{-i \restriction h} \rangle_v = g_n \varrho$, and

$$\lambda_i(h g_n \varrho) \leq \lambda_i(h g_n \varrho') < \lambda_i(h\rho)$$

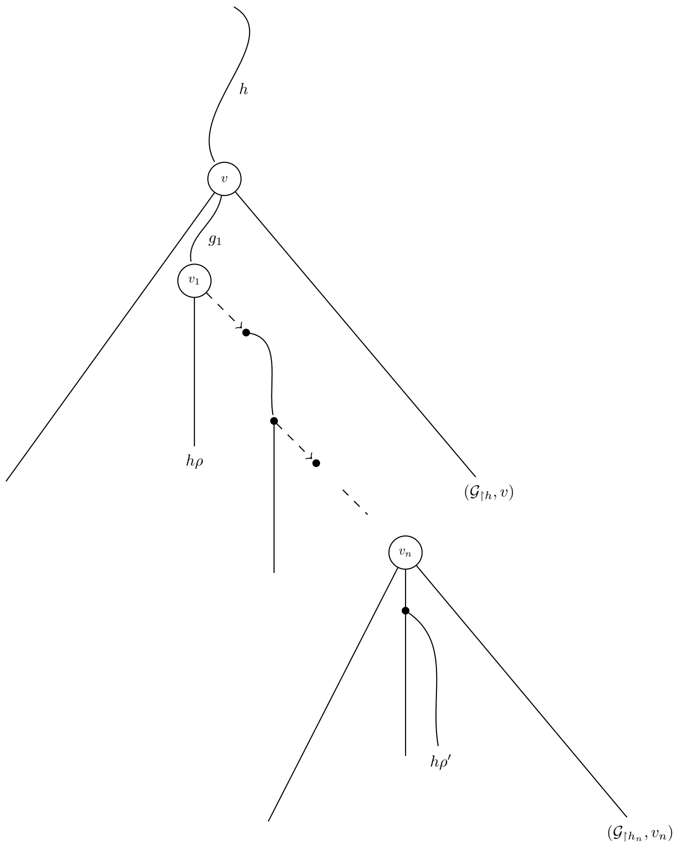

Figure 4 A strategy σ'_i with a minimum number n of deviation steps.

by (1), (2), and $g_n\varrho' = \rho'$. It follows that τ'_i is a finitely deviating strategy that is profitable for player i in $(\mathcal{G}_{\restriction h}, v)$, with less deviation steps than σ'_i, a contradiction. This completes the proof of the first statement of Proposition 13. For the second statement, it is enough to consider Example 12. ◀

Under the next hypotheses on the game or the costs, the equivalence between SPE, weak SPE, and very weak SPE holds. The first case, when the cost functions are continuous, is a classical result in game theory, see for instance [7]; the second case appears as a part of the proof of Kuhn's theorem [9].

▶ **Proposition 14.** *Let $(\mathcal{G}, v_0)$ be an initialized game, and $\bar\sigma$ be a strategy profile.*

- *If all cost functions λ_i are continuous, or even upper-semicontinuous[5], then $\bar\sigma$ is an SPE iff $\bar\sigma$ is a weak SPE iff $\bar\sigma$ is a very weak SPE.*
- *If $\mathcal{G}$ is a finite tree[6], then $\bar\sigma$ is an SPE iff $\bar\sigma$ is a weak SPE iff $\bar\sigma$ is a very weak SPE.*

Proof. We only prove the first statement for cost functions λ_i that are upper-semicontinuous. Let $\bar\sigma$ be a strategy profile in an initialized game $(\mathcal{G}, v_0)$. By Proposition 13, it remains to prove that if $\bar\sigma$ is a weak SPE, then it is an SPE, i.e., for each subgame $(\mathcal{G}_{\restriction h}, v)$, the strategy profile $\bar\sigma_{\restriction h}$ is a NE. Let σ'_i be a player i strategy in $(\mathcal{G}_{\restriction h}, v)$. If σ'_i is finitely deviating, then it is not a profitable deviation for player i w.r.t. $\bar\sigma_{\restriction h}$ by hypothesis. Therefore, suppose that

$$\rho' = \langle \sigma'_i, \sigma_{-i\restriction h}\rangle_v = g_1 g_2 \dots g_n \dots$$

such that σ'_i has a h_n-deviation step from $\bar\sigma_{\restriction h}$ for all $n \geq 1$, where $h_0 = \epsilon$ and $h_n = h_{n-1}g_n$. For each n, we define a finitely deviating strategy τ^n_i such that its deviation steps are the first n deviation steps of σ'_i, that is, τ^n_i and σ'_i are equal except in the subgame $(\mathcal{G}_{\restriction h_n}, \mathrm{First}(g_{n+1}))$ where $\tau^n_{i\ \restriction h_n} = \sigma_{i\restriction hh_n}$. By hypothesis, τ^n_i is not a profitable deviation, and thus for $\rho_n = \langle \tau^n_i, \sigma_{-i\restriction h}\rangle_v$ we have

$$\lambda_i(h \cdot \langle \bar\sigma_{\restriction h}\rangle_v) \leq \lambda_i(h\rho_n). \tag{3}$$

As $\lim_{n\to\infty} h\rho_n = h\rho'$ and λ_i is upper-semicontinuous, we get $\limsup_{n\to\infty} \lambda_i(h\rho_n) \leq \lambda_i(h\rho')$. Therefore by (3), $\lambda_i(h\cdot\langle\bar\sigma_{\restriction h}\rangle_v) \leq \lambda_i(h\rho')$ showing that σ'_i is not a profitable deviation for player i w.r.t. $\bar\sigma_{\restriction h}$. ◀

Recall that discounted sum games and quantitative reachability games are continuous. Thus for these games, the three notions of SPE, weak SPE and very weak SPE, are equivalent.

▶ **Corollary 15.** *Let $(\mathcal{G}, v_0)$ be an initialized quantitative reachability game, and $\bar\sigma$ be a strategy profile. Then $\bar\sigma$ is an SPE iff $\bar\sigma$ is a weak SPE iff $\bar\sigma$ is a very weak SPE.*

On the opposite, the initialized game of Figure 3 has a weak SPE but no SPE. Its cost function λ_2 is not upper-semicontinuous as $\lim_{n\to\infty}(v_0v_1)^n v_3^\omega = (v_0v_1)^\omega$ and $\lim_{n\to\infty}\lambda_2((v_0v_1)^n v_3^\omega)$ $= 1 > 0 = \lambda_2((v_0v_1)^\omega)$.

[5] In games where the players receive a payoff that they want to maximize, the hypothesis of upper-semicontinuity has to be replaced by lower-semicontinuity.

[6] In a finite tree game, the plays are finite sequences of vertices ending in a leaf and their cost is associated with the ending leaf. An example of such a game is depicted in Figure 2.

3 Folk Theorems

3.1 Folk Theorem for Weak SPEs

In this section, we characterize in the form of a Folk Theorem the set of all outcomes of weak SPEs. Our approach is inspired[7] by work [5] where a Folk Theorem is given for the set of outcomes of SPEs in games with cost functions that are upper-semicontinuous and have finite range. In this aim we define a nonincreasing sequence of sets of plays that initially contain all the plays, and then loose, step by step, some plays that for sure are not outcomes of a weak SPE, until finally reaching a fixpoint.

Let $(\mathcal{G}, v_0)$ be a game. For an ordinal α and a history $hv \in \mathrm{Hist}(v_0)$, let us consider the set $\mathbf{P}_\alpha(hv) = \{\rho \mid \rho \text{ is a } \textit{potential} \text{ outcome of a weak NE in } (\mathcal{G}_{\restriction h}, v) \text{ at step } \alpha\}$. This set is defined by induction on α as follows:

▶ **Definition 16.** Let $(\mathcal{G}, v_0)$ be a quantitative game. The set $\mathbf{P}_\alpha(hv)$ is defined as follows for each ordinal α and history $hv \in \mathrm{Hist}(v_0)$:

- For $\alpha = 0$,
$$\mathbf{P}_\alpha(hv) = \{\rho \mid \rho \text{ is a play in } (\mathcal{G}_{\restriction h}, v)\}. \tag{4}$$
- For a successor ordinal $\alpha + 1$,
$$\mathbf{P}_{\alpha+1}(hv) = \mathbf{P}_\alpha(hv) \setminus \mathbf{E}_\alpha(hv) \tag{5}$$
such that $\rho \in \mathbf{E}_\alpha(hv)$ (see Figure 5) iff
  - there exists a history h', $hv \leq h' < h\rho$, and $\mathrm{Last}(h') \in V_i$ for some i,
  - there exists a vertex v', $h'v' \not< h\rho$,
  - such that $\forall \rho' \in \mathbf{P}_\alpha(h'v')$: $\lambda_i(h\rho) > \lambda_i(h'\rho')$.
- For a limit ordinal α:
$$\mathbf{P}_\alpha(hv) = \bigcap_{\beta<\alpha} \mathbf{P}_\beta(hv). \tag{6}$$

Notice that an element ρ of $\mathbf{P}_\alpha(hv)$ is a play in $(\mathcal{G}_{\restriction h}, v)$ (and not in $(\mathcal{G}, v_0)$). Therefore it starts with vertex v, and $h\rho$ is a play in $(\mathcal{G}, v_0)$. For $\alpha + 1$ being a successor ordinal, play $\rho \in \mathbf{E}_\alpha(hv)$ is erased from $\mathbf{P}_\alpha(hv)$ because for all $\rho' \in \mathbf{P}_\alpha(h'v')$, player i pays a lower cost $\lambda_i(h'\rho') < \lambda_i(h\rho)$, which means that ρ is no longer a potentiel outcome of a weak NE in $(\mathcal{G}_{\restriction h}, v)$.

The sequence $(\mathbf{P}_\alpha(hv))_\alpha$ is nonincreasing by definition, and reaches a fixpoint in the following sense.

▶ **Proposition 17.** *There exists an ordinal α_* such that $\mathbf{P}_{\alpha_*}(hv) = \mathbf{P}_{\alpha_*+1}(hv)$ for all histories $hv \in \mathrm{Hist}(v_0)$.*

Proof. Let us fix a history $hv \in \mathrm{Hist}(v_0)$. The sequence $(\mathbf{P}_\alpha(hv))_\alpha$ reaches a fixpoint as soon as there exists α such that $\mathbf{P}_\alpha(hv) = \mathbf{P}_{\alpha+1}(hv)$. Indeed it follows that $\mathbf{P}_{\alpha+1}(hv) = \mathbf{P}_{\alpha+2}(hv)$ and then $\mathbf{P}_\alpha(hv) = \mathbf{P}_\beta(hv)$ for all $\beta > \alpha$. As the sequence $(\mathbf{P}_\alpha(hv))_\alpha$ is nonincreasing, this happens at the latest with α being equal to the cardinality of $\mathbf{P}_0(hv)$. Therefore with $\alpha_* = |V^\omega|$ being an ordinal greater than or equal to the cardinality of the set of all plays of $\mathcal{G}$, we get $\mathbf{P}_{\alpha_*}(hv) = \mathbf{P}_{\alpha_*+1}(hv)$ for all $hv \in \mathrm{Hist}(v_0)$. ◀

In the sequel, α_* always refers to the ordinal mentioned in Proposition 17.

Our Folk Theorem for weak SPEs is the next one.

---

[7] Our approach is however different.

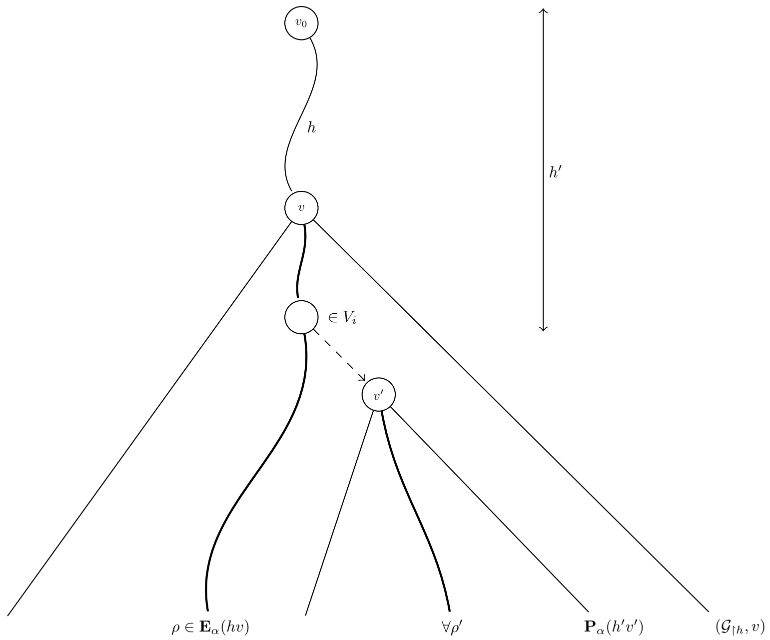

Figure 5 $\rho \in \mathbf{E}_\alpha(hv)$.

▶ **Theorem 18.** *Let $(\mathcal{G}, v_0)$ be a quantitative game. There exists a weak SPE in $(\mathcal{G}, v_0)$ with outcome ρ iff $\mathbf{P}_{\alpha_*}(hv) \neq \emptyset$ for all $hv \in \mathrm{Hist}(v_0)$, and $\rho \in \mathbf{P}_{\alpha_*}(v_0)$.*

Before proving this theorem, we illustrate it with an example.

▶ **Example 19.** Consider the example of Figure 3. Clearly, as $\mathbf{P}_0(hv_2)$ only contains the play v_2^ω, then $\mathbf{P}_\alpha(hv_2) = \{v_2^\omega\}$ for all α. Similarly $\mathbf{P}_\alpha(hv_3) = \{v_3^\omega\}$ for all α. Let us detail the computation of $\mathbf{P}_\alpha(hv_0)$ and $\mathbf{P}_\alpha(hv_1)$.

- $\alpha = 0$. For history hv_0, we have $\rho = (v_0v_1)^\omega \in \mathbf{E}_0(hv_0)$, since $\lambda_1(h\rho) > \lambda_1(hv_0\rho')$ where $\rho' = v_2^\omega$ is the unique play of $\mathbf{P}_0(hv_0v_2)$. Similarly, for all $n \geq 1$, we have $\rho = (v_0v_1)^n v_0 v_2^\omega \in \mathbf{E}_0(hv_0)$, since $\lambda_2(h\rho) > \lambda_2(hv_0v_1\rho')$ where $\rho' = v_3^\omega$ is the unique play of $\mathbf{P}_0(hv_0v_1v_3)$. Thus $\mathbf{E}_0(hv_0) = \{(v_0v_1)^\omega\} \cup (v_0v_1)^+ v_0 v_2^\omega$ and

$$\mathbf{P}_1(hv_0) = \{v_0 v_2^\omega\} \cup (v_0v_1)^+ v_3^\omega.$$

  For history hv_1, with the same kind of computations, we get $\mathbf{E}_0(hv_1) = \{(v_1v_0)^\omega\} \cup (v_1v_0)^+ v_2^\omega$ and

$$\mathbf{P}_1(hv_1) = v_1(v_0v_1)^* v_3^\omega.$$

- $\alpha = 1$. For history hv_0, we have $\rho = v_0 v_2^\omega \in \mathbf{E}_1(hv_0)$. Indeed $\lambda_1(h\rho) > \lambda_1(hv_0\rho')$ for all $\rho' \in \mathbf{P}_1(hv_0v_1) = v_1(v_0v_1)^* v_3^\omega$. Notice that at the previous step, $\rho = v_0 v_2^\omega \notin \mathbf{E}_0(hv_0)$. Indeed $\mathbf{P}_1(hv_0v_1) \subsetneq \mathbf{P}_0(hv_0v_1)$, and $\lambda_1(h\rho) \leq \lambda_1(hv_0\rho')$ for $\rho' = (v_1v_0)^\omega \in \mathbf{P}_0(hv_0v_1)$.[8] Play $v_0 v_2^\omega$ is the only one that is removed from $\mathbf{P}_1(hv_0)$, and no play can be removed from $\mathbf{P}_1(hv_1)$. Therefore:

$$\mathbf{P}_2(hv_0) = (v_0v_1)^+ v_3^\omega, \quad \mathbf{P}_2(hv_1) = v_1(v_0v_1)^* v_3^\omega.$$

[8] This shows that as the sequence $(\mathbf{P}_\alpha(hv))_\alpha$ is nonincreasing for each hv, a play that is not removed from some $\mathbf{P}_\alpha(hv)$ can be removed later from some $\mathbf{P}_\beta(hv)$ with $\beta > \alpha$.

- $\alpha = 2$. One checks that $\mathbf{P}_3(hv_0) = \mathbf{P}_2(hv_0)$, and $\mathbf{P}_3(hv_1) = \mathbf{P}_2(hv_1)$. Hence the fixpoint is reached with $\alpha_* = 2$, with $\mathbf{P}_{\alpha_*}(hv_0) = (v_0v_1)^+v_3^\omega$, $\mathbf{P}_{\alpha_*}(hv_1) = v_1(v_0v_1)^*v_3^\omega$, $\mathbf{P}_{\alpha_*}(hv_2) = \{v_2^\omega\}$, and $\mathbf{P}_{\alpha_*}(hv_3) = \{v_3^\omega\}$. Therefore, the set of outcomes of weak SPEs in this game is equal to $(v_0v_1)^+v_3^\omega$. The weak SPE depicted in Figure 3 has outcome $v_0v_1v_3^\omega$.

The proof of Theorem 18 follows from Lemmas 20 and 21.

▶ **Lemma 20.** *If $(\mathcal{G}, v_0)$ has a weak SPE $\bar\sigma$, then $\mathbf{P}_{\alpha_*}(hv) \neq \emptyset$ for all $hv \in \text{Hist}(v_0)$, and $\langle\bar\sigma\rangle_{v_0} \in \mathbf{P}_{\alpha_*}(v_0)$.*

Proof. Let us show, by induction on α, that $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}_\alpha(hv)$ for all $hv \in \text{Hist}(v_0)$.

For $\alpha = 0$, we have $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}_\alpha(hv)$ by definition of $\mathbf{P}_0(hv)$.

Let $\alpha + 1$ be a successor ordinal. By induction hypothesis, we have that $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}_\alpha(hv)$ for all $hv \in \text{Hist}(v_0)$. Suppose that there exists hv such that $\langle\bar\sigma_{\restriction h}\rangle_v \notin \mathbf{P}_{\alpha+1}(hv)$, i.e. $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{E}_\alpha(hv)$. This means that there is a history $h' = hg \in \text{Hist}_i$ for some $i \in \Pi$ with $hv \leq h' < h\rho$, and there exists a vertex v' with $h'v' \not< h\rho$, such that $\forall\rho' \in \mathbf{P}_\alpha(h'v')$, $\lambda_i(h\cdot\langle\bar\sigma_{\restriction h}\rangle_v) > \lambda_i(h'\rho')$. In particular, by induction hypothesis

$$\lambda_i(h\cdot\langle\bar\sigma_{\restriction h}\rangle_v) > \lambda_i(h'\cdot\langle\bar\sigma_{\restriction h'}\rangle_{v'}). \tag{7}$$

Let us consider the player i strategy σ'_i in $(\mathcal{G}_{\restriction h}, v)$ such that $g\cdot\langle\bar\sigma_{\restriction h'}\rangle_{v'}$ is consistent with σ'_i. Then σ'_i is a finitely deviating strategy with the (unique) g-deviation step from $\bar\sigma_{\restriction h}$. This strategy is a profitable deviation for player i in $(\mathcal{G}_{\restriction h}, v)$ by (7), a contradiction with $\bar\sigma$ being a weak SPE.

Let α be a limit ordinal. By induction hypothesis $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}_\beta(hv), \forall\beta < \alpha$. Therefore $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}_\alpha(hv) = \bigcap_{\beta<\alpha}\mathbf{P}_\beta(hv)$. ◀

▶ **Lemma 21.** *Suppose that $\mathbf{P}_{\alpha_*}(hv) \neq \emptyset$ for all $hv \in \text{Hist}(v_0)$, and let $\rho \in \mathbf{P}_{\alpha_*}(v_0)$. Then $(\mathcal{G}, v_0)$ has a weak SPE with outcome ρ.*

Proof. We are going to show how to construct a very weak SPE $\bar\sigma$ (and thus a weak SPE by Proposition 13) with outcome ρ. The construction of $\bar\sigma$ is done step by step thanks to a progressive labeling of the histories $hv \in \text{Hist}(v_0)$. Let us give an intuitive idea of the construction of $\bar\sigma$. Initially, we partially construct $\bar\sigma$ such that it produces an outcome in $(\mathcal{G}, v_0)$ equal to $\rho \in \mathbf{P}_{\alpha_*}(v_0)$; we also label each non-empty prefix of ρ by ρ. Then we consider a shortest non-labeled history $h'v'$, and we correctly choose some $\rho' \in \mathbf{P}_{\alpha_*}(h'v')$ (we will see later how). We continue the construction of $\bar\sigma$ such that it produces the outcome ρ' in $(\mathcal{G}_{\restriction h'}, v')$, and for each non-empty prefix g of ρ', we label $h'g$ by ρ' (notice that the prefixes of h' have already been labeled by choice of h'). And so on. In this way, the labeling is a map $\gamma : \text{Hist}(v_0) \to \bigcup_{hv} \mathbf{P}_{\alpha_*}(hv)$ that allows to recover from $h'g$ the outcome ρ' of $\bar\sigma_{\restriction h'}$ in $(\mathcal{G}_{\restriction h'}, v')$ of which g is prefix. Let us now go into the details.

Initially, none of the histories is labeled. We start with history v_0 and the given play $\rho \in \mathbf{P}_{\alpha_*}(v_0)$. The strategy profile $\bar\sigma$ is partially defined such that $\langle\bar\sigma\rangle_{v_0} = \rho$, that is, if $\rho = \rho_0\rho_1\ldots$, then $\sigma_i(\rho_{\leq n}) = \rho_{n+1}$ for all $\rho_n \in V_i$ and $i \in \Pi$. The non-empty prefixes h of ρ are all labeled with $\gamma(h) = \rho$.

At the following steps, we consider a history $h'v'$ that is not yet labeled, but such that h' has already been labeled. By induction, $\gamma(h') = \langle\bar\sigma_{\restriction h}\rangle_v$ such that $hv \leq h'$. Suppose that $\text{Last}(h') \in V_i$, we then choose a play $\rho' \in \mathbf{P}_{\alpha_*}(h'v')$ such that (see Figure 6)

$$\lambda_i(h\cdot\langle\bar\sigma_{\restriction h}\rangle_v) \leq \lambda_i(h'\rho'). \tag{8}$$

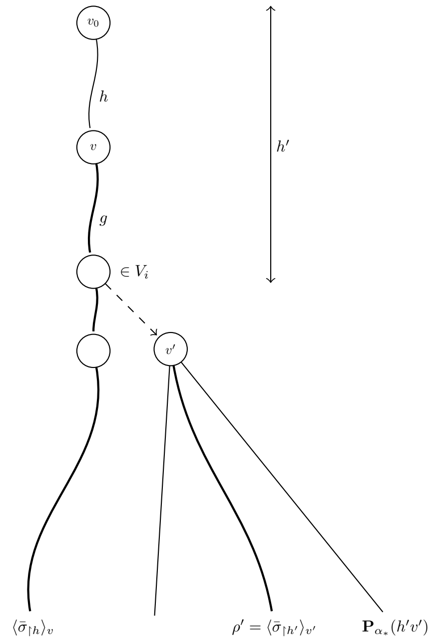

Figure 6 Construction of a very weak SPE $\bar\sigma$.

Such a play ρ' exists for the following reasons. By induction, we know that $\langle \bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}_{\alpha_*}(hv)$. Since $\mathbf{P}_{\alpha_*}(hv) = \mathbf{P}_{\alpha_*+1}(hv)$ by Proposition 17, we have $\langle \bar\sigma_{\restriction h}\rangle_v \notin \mathbf{E}_{\alpha_*}(hv)$, and we get the existence of ρ' by definition of $\mathbf{E}_{\alpha_*}(hv)$. We continue to construct $\bar\sigma$ such that $\langle \bar\sigma_{\restriction h'}\rangle_{v'} = \rho'$, i.e. if $\rho' = \rho'_0\rho'_1\ldots$, then $\sigma_i(h'\rho'_{\leq n}) = h'\rho'_{n+1}$ for all $\rho'_n \in V_i$ and $i \in \Pi$. For all non-empty prefixes g of ρ', we define $\gamma(h'g) = \rho'$ (notice that the prefixes of h' have already been labeled).

Let us show that the constructed profile $\bar\sigma$ is a very weak SPE. Consider a history $hv \in \mathrm{Hist}_i$ for some $i \in \Pi$, and a one-shot deviating strategy σ'_i from $\bar\sigma_{\restriction h}$ in the subgame $(\mathcal{G}_{\restriction h}, v)$. Let v' be such that $\sigma'_i(v) = v'$. By definition of $\bar\sigma$, we have $\gamma(hv) = \langle \bar\sigma_{\restriction g}\rangle_u$ for some history $gu \leq hv$ and $h\cdot\langle \bar\sigma_{\restriction h}\rangle_v = g\cdot\langle \bar\sigma_{\restriction g}\rangle_u$; and we have also $\gamma(hvv') = \langle \bar\sigma_{\restriction hv}\rangle_{v'}$. Moreover $\lambda_i(g\cdot\langle \bar\sigma_{\restriction g}\rangle_u) \leq \lambda_i(hv\cdot\langle \bar\sigma_{\restriction hv}\rangle_{v'})$ by (8), and $\lambda_i(hv\cdot\langle \bar\sigma_{\restriction hv}\rangle_{v'}) = \lambda_i(h\cdot\langle \sigma'_i, \sigma_{-i\restriction h}\rangle_v)$ because σ'_i is one-shot deviating. Therefore

$$\lambda_i(h\cdot\langle \bar\sigma_{\restriction h}\rangle_v) = \lambda_i(g\cdot\langle \bar\sigma_{\restriction g}\rangle_u) \leq \lambda_i(hv\cdot\langle \bar\sigma_{\restriction hv}\rangle_{v'}) = \lambda_i(h\cdot\langle \sigma'_i, \sigma_{-i\restriction h}\rangle_v)$$

which shows that $\bar\sigma_{\restriction h}$ is a very weak NE in $(\mathcal{G}_{\restriction h}, v)$. Hence $\bar\sigma$ is a very weak SPE, and thus also a weak SPE. ◀

The next lemma will be useful in Sections 4 and 5. It states that if a play ρ belongs to $\mathbf{P}_\alpha(hv)$, then each of its suffixes ρ_1 also belongs to $\mathbf{P}_\alpha(hh_1v_1)$ such that $h_1\rho_1 = \rho$ and $v_1 = \mathrm{First}(\rho_1)$.

▶ **Lemma 22.** *Let $\rho \in \mathbf{P}_\alpha(hv)$. Then for all $h_1\rho_1 = \rho$, we have $\rho_1 \in \mathbf{P}_\alpha(hh_1v_1)$ with $v_1 = \mathrm{First}(\rho_1)$.*

Proof. The proof is by induction on α. The lemma trivially holds for $\alpha = 0$ by definition of $\mathbf{P}_0(hv)$.

Let $\alpha + 1$ be a successor ordinal. Let $\rho \in \mathbf{P}_{\alpha+1}(hv)$ and $h_1\rho_1 = \rho$ with $v_1 = \text{First}(\rho_1)$. As $\mathbf{P}_{\alpha+1}(hv) \subseteq \mathbf{P}_\alpha(hv)$, by induction hypothesis, we have $\rho_1 \in \mathbf{P}_\alpha(hh_1v_1)$. Suppose that $\rho_1 \in \mathbf{E}_\alpha(hh_1v_1)$ (hence using a history h' and a vertex v' as in Definition 16). Then one can easily check by definition of $\mathbf{E}_\alpha(hh_1v_1)$ that $\rho \in \mathbf{E}_\alpha(hv)$ (by using the same h' and v'), which is a contradiction with $\rho \in \mathbf{P}_{\alpha+1}(hv)$. Therefore $\rho_1 \in \mathbf{P}_\alpha(hh_1v_1) \setminus \mathbf{E}_\alpha(hh_1v_1) = \mathbf{P}_{\alpha+1}(hh_1v_1)$.

Let α be a limit ordinal, and suppose that $\rho \in \mathbf{P}_\alpha(hv)$. As $\rho \in \mathbf{P}_\beta(hv)$ for all $\beta < \alpha$, we have $\rho_1 \in \mathbf{P}_\beta(hh_1v_1)$ by induction hypothesis. It follows that $\rho_1 \in \mathbf{P}_\alpha(hh_1v_1) = \bigcap_{\beta<\alpha} \mathbf{P}_\beta(hh_1v_1)$. ◀

3.2 Folk Theorem for SPEs

In this section, as for weak SPEs, we characterize in the form of a Folk Theorem the set of all outcomes of SPEs. Nevertheless, we here need a more complex characterization with adapted sets $\mathbf{P}'_\alpha(hv)$, and this characterization only holds for cost functions that are upper-semicontinuous.[9] The main difference appears in the definition of sets $\mathbf{E}'_\alpha(hv)$ that will be used in place of $\mathbf{E}_\alpha(hv)$. Indeed we will see that the set $\mathbf{P}_\alpha(h'v')$ of Figure 5 has to be replaced by a more complex set $\mathbf{D}^{H,i}_\alpha(h'v')$.

Let $(\mathcal{G}, v_0)$ be a game. For an ordinal α and a history $hv \in \text{Hist}(v_0)$, as in the previous section, we consider the set $\mathbf{P}'_\alpha(hv) = \{\rho \mid \rho$ is a potential outcome of a NE[10] in $(\mathcal{G}_{\restriction h}, v)$ at step $\alpha\}$. In order to define these sets $\mathbf{P}'_\alpha(hv)$, we need to introduce new definitions. In Definition 4, we have introduced the notion of deviation step of a strategy from a given strategy profile. We here propose another concept of deviation step in relation with two plays (see Figure 7).

▶ **Definition 23.** Let $h'v' \in \text{Hist}(v_0)$ and $\rho' \in (\mathcal{G}_{\restriction h'}, v')$. Let $h_1u_1, h_2u_2 \in \text{Hist}(v_0)$ with $u_1, u_2 \in V$, and ρ_1 in $\mathbf{P}'_\alpha(h_1u_1)$. We say that ρ' has a h_2u_2*-deviation step* from ρ_1 if $h'v' \leq h_1u_1 < h_2u_2 < h'\rho'$ and $h'\rho' \wedge h_1\rho_1 = h_2$.

Let us denote by $\mathbf{H}(h'v')$ the set of all histories h_2u_2 such that $h'v' < h_2u_2$. Given a player i and $H \subseteq \mathbf{H}(hv)$, the next definition introduces the notion of (H, i)-decomposition of a play ρ'. Such a play has a finite or infinite number of deviation steps such that for each h_nu_n-deviation step, the history h_nu_n belongs to H. Figure 8 illustrates the second case of Definition 24, with the deviation steps highlighted with dashed edges. This definition also introduces the set $\mathbf{D}^{H,i}_\alpha(h'v')$ composed of plays with a maximal (H, i)-decomposition.

▶ **Definition 24.** Let $h'v' \in \text{Hist}(v_0)$ and $\rho' \in (\mathcal{G}_{\restriction h'}, v')$. Let i be a player, and $H \subseteq \mathbf{H}(h'v')$.

- ρ' has an *infinite* (H, i)*-decomposition* $\rho' = g_1g_2 \dots g_n \dots$ if
  - for all $n \geq 1$, $\text{Last}(g_n) \in V_i$
  - for all $n \geq 2$, $h_nu_n \in H$, and ρ' has a h_nu_n-deviation step from some $\rho_{n-1} \in \mathbf{P}'_\alpha(h_{n-1}u_{n-1})$

  where $h_1 = h'$, $h_{n+1} = h_ng_n$, and $u_n = \text{First}(g_n)$ $\forall n \geq 1$.
- ρ' has a *finite* (H, i)*-decomposition* $\rho' = g_1g_2 \dots g_m\varrho'$ if
  - for all n, $1 \leq n \leq m$, $\text{Last}(g_n) \in V_i$
  - for all n, $2 \leq n \leq m+1$, $h_nu_n \in H$, and ρ' has a h_nu_n-deviation step from some $\rho_{n-1} \in \mathbf{P}'_\alpha(h_{n-1}u_{n-1})$
  - $\varrho' \in \mathbf{P}'_\alpha(h_{m+1}u_{m+1})$

[9] For games where players receive a payoff that they want to maximize, a similar Folk Theorem also exists for lower-semicontinuous cost functions.

[10] instead of a weak NE

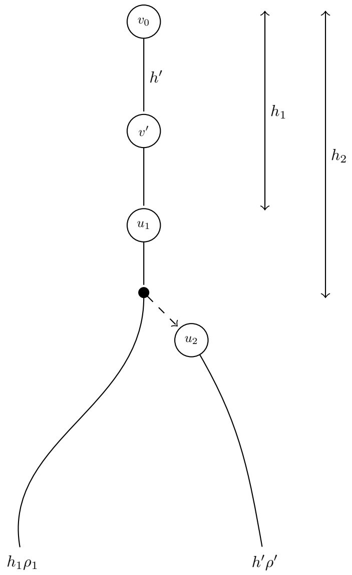

Figure 7 Deviation step.

where $h_1 = h'$, $h_{n+1} = h_n g_n$, $u_n = \text{First}(g_n)$ $\forall n, 1 \leq n \leq m$, and $u_{m+1} = \text{First}(\varrho')$.

We denote by $\mathbf{D}^{H,i}_{\alpha}(h'v')$ the set of plays ρ' with a *maximal* (H,i)*-decomposition* in the following sense:

- ρ' has an infinite (H,i)-decomposition, or
- ρ' has a finite (H,i)-decomposition $\rho' = g_1 g_2 \dots g_m \varrho'$, and there exists no ρ'' with a (H,i)-decomposition $\rho'' = g_1 g_2 \dots g_{m'} \dots$ or $\rho'' = g_1 g_2 \dots g_{m'} \varrho''$ such that $\rho' \frown \rho'' = g_1 g_2 \dots g_{m+1}$.

In the previous definition, H can be chosen finite, infinite, or empty. If $H = \emptyset$, then for all $i \in \Pi$, for each $\rho' \in \mathbf{D}^{H,i}_{\alpha}(h'v')$, ρ' has no deviation step and thus $\rho' \in \mathbf{P}'_{\alpha}(h'v')$. This means that $\mathbf{D}^{H,i}_{\alpha}(h'v') = \mathbf{P}'_{\alpha}(h'v')$ in this case.

We are ready to define the sets $\mathbf{P}'_{\alpha}(hv)$ by induction on α. The definition is similar to the one of $\mathbf{P}_{\alpha}(hv)$, except that when we erase $\rho \in \mathbf{E}'_{\alpha}(hv)$ from $\mathbf{P}'_{\alpha}(hv)$, we use some set $\mathbf{D}^{H,i}_{\alpha}(h'v')$ in place of $\mathbf{P}'_{\alpha}(h'v')$:

▶ **Definition 25.** Let $(\mathcal{G}, v_0)$ be a quantitative game. The set $\mathbf{P}'_{\alpha}(hv)$ is defined as follows for each ordinal α and history $hv \in \text{Hist}(v_0)$:

- For $\alpha = 0$,
$$\mathbf{P}'_{\alpha}(hv) = \{\rho \mid \rho \text{ is a play in } (\mathcal{G}_{\upharpoonright h}, v)\}. \tag{9}$$
- For a successor ordinal $\alpha + 1$,
$$\mathbf{P}'_{\alpha+1}(hv) = \mathbf{P}'_{\alpha}(hv) \setminus \mathbf{E}'_{\alpha}(hv) \tag{10}$$
such that $\rho \in \mathbf{E}'_{\alpha}(hv)$ iff
  - there exists a history h', $hv \leq h' < h\rho$, and $\text{Last}(h') \in V_i$ for some i,
  - there exists a vertex v', $h'v' \not< h\rho$,
  - there exists $H \subseteq \mathbf{H}(h'v')$,
  - such that $\forall \rho' \in \mathbf{D}^{H,i}_{\alpha}(h'v')$: $\lambda_i(h\rho) > \lambda_i(h'\rho')$.

  (see Figure 5 where $\mathbf{E}_{\alpha}(hv)$ is replaced by $\mathbf{E}'_{\alpha}(hv)$, and $\mathbf{P}'_{\alpha}(h'v')$ is replaced by $\mathbf{D}^{H,i}_{\alpha}(h'v')$).

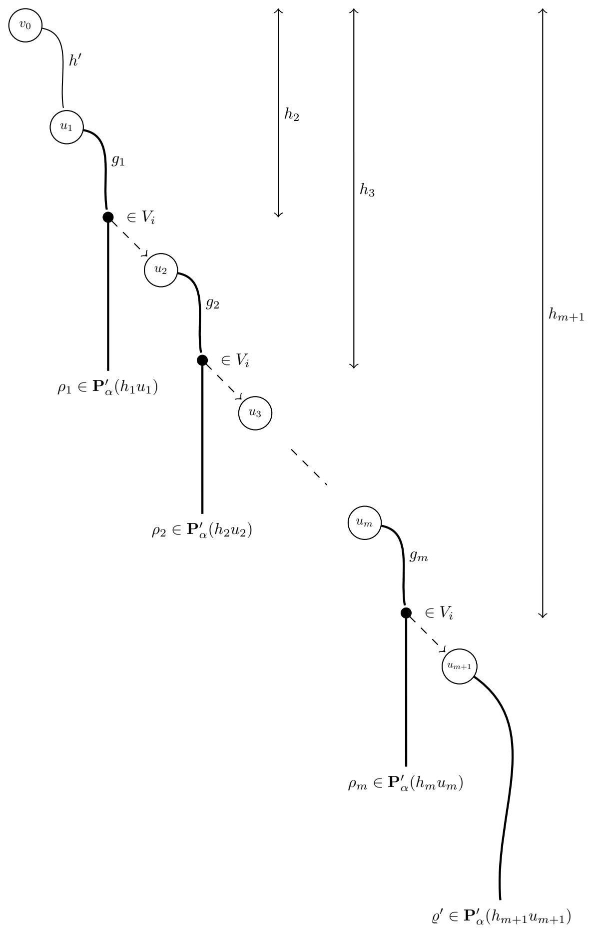

Figure 8 $\rho' = g_1g_2 \dots g_m\varrho'$ with a finite (H, i)-decomposition.

- For a limit ordinal α:
$$\mathbf{P}'_\alpha(hv) = \bigcap_{\beta<\alpha} \mathbf{P}'_\beta(hv). \tag{11}$$

Let us comment the case $\alpha + 1$. When $H = \emptyset$, we have $\mathbf{D}^{H,i}_\alpha(h'v') = \mathbf{P}'_\alpha(h'v')$. Hence we recover the previous situation of weak SPEs. Using different sets $H \subseteq \mathbf{H}(h'v')$ and $\mathbf{D}^{H,i}_\alpha(h'v')$ allow to have sets $\mathbf{E}'_\alpha(hv)$ bigger than $\mathbf{E}_\alpha(hv)$, and thus more plays removed from $\mathbf{P}'_\alpha(hv)$ than in $\mathbf{P}_\alpha(hv)$. This situation will be illustrated in Example 28 hereafter.

The sequence $(\mathbf{P}'_\alpha(hv))_\alpha$ is nonincreasing by definition, and reaches a fixpoint in the following sense (the proof is the same as for Proposition 17).

▶ **Proposition 26.** *There exists an ordinal β_* such that $\mathbf{P}'_{\beta_*}(hv) = \mathbf{P}'_{\beta_*+1}(hv)$ for all histories $hv \in \mathrm{Hist}(v_0)$.*

Our Folk Theorem for SPEs is the next one. The second statement requires to work with upper-semicontinuous cost functions λ_i, $i \in \Pi$.

▶ **Theorem 27.** *Let $(\mathcal{G}, v_0)$ be a quantitative game.*

- *If there exists an SPE in $(\mathcal{G}, v_0)$ with outcome ρ, then $\mathbf{P}'_{\beta_*}(hv) \neq \emptyset$ for all $hv \in \mathrm{Hist}(v_0)$, and $\rho \in \mathbf{P}'_{\beta_*}(v_0)$.*

- *Suppose that all cost functions λ_i are upper-semicontinuous. If $\mathbf{P}'_{\beta_*}(hv) \neq \emptyset$ for all $hv \in$ Hist(v_0), then there exists an SPE in $(\mathcal{G}, v_0)$ with outcome ρ, for all $\rho \in \mathbf{P}'_{\beta_*}(v_0)$.*

▶ **Example 28.** Before proving this theorem, we illustrate it with the example of Figure 3 by showing that this game has no SPE (as stated in [16]). Let us compute the sets $\mathbf{P}'_\alpha(hv)$ and let us show that $\mathbf{P}'_3(hv_2) = \emptyset$. By the first statement of Theorem 27, we will get that there is no SPE.

We have to do computations that are more complex than the ones of Example 19, due to the usage of sets $\mathbf{D}^{H,i}_\alpha(h'v')$, with $H \subseteq \mathbf{H}(h'v')$, instead of $\mathbf{P}_\alpha(h'v')$. Clearly, by definition $\mathbf{P}'_0(hv) = \mathbf{P}_0(hv)$ for all hv, and we have $\mathbf{P}'_\alpha(hv_2) = \{v_2^\omega\}$ and $\mathbf{P}'_\alpha(hv_3) = \{v_3^\omega\}$ for all α as in Example 19.

Let us first illustrate Definition 24 with the computation of $\mathbf{D}^{H,2}_0(v_0)$ with $H = \{v_0v_1v_0, (v_0v_1)^2v_0\}$. The play $\rho'_1 = (v_0v_1)^2v_0v_2^\omega$ has a $(H,2)$-decomposition $g_1g_2\varrho'_1$ with $g_1 = g_2 = v_0v_1$ and $\varrho' = v_2^\omega$ (two deviation steps). Indeed ρ'_1 has a $v_0v_1v_0$-deviation step from $v_0v_1v_3^\omega \in \mathbf{P}'_\alpha(v_0)$ and a $(v_0v_1)^2v_0$-deviation step from $(v_0v_1)^2v_3^\omega \in \mathbf{P}'_\alpha(v_0v_1v_0)$, $\mathrm{Last}(g_1), \mathrm{Last}(g_2) \in V_2$, and $\varrho'_1 \in \mathbf{P}'_\alpha((v_0v_1)^2v_0)$. This play ρ'_1 belongs to $\mathbf{D}^{H,2}_0(v_0)$ because its $(H,2)$-decomposition uses the two possible steps of H. On the contrary $\rho'_2 = (v_0v_1)^2v_3^\omega$ has a $(H,2)$-decomposition $g_1\varrho'_2$ with $\varrho'_2 = v_0v_1v_3^\omega$, and does not belong to $\mathbf{D}^{H,2}_0(v_0)$ because $\rho'_1{}^\frown\rho'_2 = g_1g_2$. One can check that $\mathbf{D}^{H,2}(v_0) = \{v_0v_2^\omega, v_0v_1v_0v_2^\omega\} \cup \{(v_0v_1)^2\varrho' \mid \varrho' \in (\mathcal{G}_{\restriction (v_0v_1)^2}, v_0)\}$.

We can now detail the computation of $\mathbf{P}'_\alpha(hv_0)$ and $\mathbf{P}'_\alpha(hv_1)$.

- $\alpha = 0$. For history hv_0, we have $\mathbf{E}'_0(hv_0) \supseteq \mathbf{E}_0(hv_0) = \{(v_0v_1)^\omega\} \cup (v_0v_1)^+v_0v_2^\omega$ because $\mathbf{D}^{H,i}_0(h'v') = \mathbf{P}'_0(h'v') = \mathbf{P}_0(h'v')$ when $H = \emptyset$. In fact, one checks that $\mathbf{E}'_0(hv_0) = \mathbf{E}_0(hv_0)$, and thus
$$\mathbf{P}'_1(hv_0) = \mathbf{P}_1(hv_0) = \{v_0v_2^\omega\} \cup (v_0v_1)^+v_3^\omega.$$
For instance, if we try to remove $v_0v_2^\omega$ with cost $(1,2)$ from $\mathbf{P}'_0(hv_0)$, we have to use some H such that $\mathbf{D}^{H,1}_0(hv_0v_1) \subseteq v_1(v_0v_1)^*v_3^\omega$ (with cost $(0,1)$). Such a H does not exist. For history hv_1, we also have $\mathbf{E}'_0(hv_1) = \mathbf{E}_0(hv_1)$ and
$$\mathbf{P}'_1(hv_1) = \mathbf{P}_1(hv_1) = v_1(v_0v_1)^*v_3^\omega.$$
- $\alpha = 1$. Again we have
$$\mathbf{P}'_2(hv_0) = \mathbf{P}_2(hv_0) = (v_0v_1)^+v_3^\omega, \quad \mathbf{P}'_2(hv_1) = \mathbf{P}_2(hv_1) = v_1(v_0v_1)^*v_3^\omega.$$
- $\alpha = 2$. A difference appears at this step: $\mathbf{P}_3(hv_1) = v_1(v_0v_1)^*v_3^\omega$ whereas
$$\mathbf{P}'_3(hv_1) = \emptyset.$$
Indeed $\mathbf{E}'_2(hv_1) = \mathbf{P}'_2(hv_1)$. Consider for instance $\rho = v_1v_3^\omega \in \mathbf{P}'_2(hv_1)$, and $H = hv_1v_0(v_1v_0)^+ \subseteq \mathbf{H}(hv_1v_0)$. Then $\mathbf{D}^{H,2}_2(hv_1v_0)$ has a unique play $\rho' = (v_0v_1)^\omega$, and $\lambda_2(h\rho) > \lambda_2(hv_1\rho')$.

Proof of Theorem 27. We begin by the first statement. Let $\bar\sigma$ be an SPE. As in the proof of Lemma 20, let us show by induction on α that $\langle \bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}'_\alpha(hv)$ for all $hv \in$ Hist(v_0).

For $\alpha = 0$, we have $\langle \bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}'_\alpha(hv)$ by definition of $\mathbf{P}'_0(hv)$.

Let $\alpha + 1$ be a successor ordinal. By induction hypothesis, we have that $\langle \bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}'_\alpha(hv)$ for all hv. Suppose that $\langle \bar\sigma_{\restriction h}\rangle_v \notin \mathbf{P}'_{\alpha+1}(hv)$, i.e. $\langle \bar\sigma_{\restriction h}\rangle_v \in \mathbf{E}'_\alpha(hv)$. This means that there is a history $h' = hg \in$ Hist$_i$ for some i with $hv \leq h' < h\rho$, there exists a vertex v' with $h'v' \not< h\rho$, and there exists $H \subseteq \mathbf{H}(h'v')$, such that $\forall \rho' \in \mathbf{D}^{H,i}_\alpha(h'v')$,
$$\lambda_i(h \cdot \langle \bar\sigma_{\restriction h}\rangle_v) > \lambda_i(h'\rho'). \qquad (12)$$

Let us consider player i strategy σ'_i in the subgame $(\mathcal{G}_{\restriction h}, v)$ such that σ'_i coincide with $\sigma_{i\restriction h}$ except that

$$\sigma'_i(h'_1) = v'_1, \text{ for all } hh'_1v'_1 \in H \cup \{h'v'\}. \tag{13}$$

Let $g\rho^* = \langle \sigma'_i, \sigma_{-i\restriction h}\rangle_v$. We get that ρ^* has a maximal (H,i)-decomposition such that each $\rho_{n-1} \in \mathbf{P}'_\alpha(h_{n-1}u_{n-1})$ of Definition 24 is equal to $\langle \bar\sigma_{\restriction h_{n-1}}\rangle_{u_{n-1}}$ (this play belongs to $\mathbf{P}'_\alpha(h_{n-1}u_{n-1})$ by induction hypothesis). Each deviation step of ρ^* in the sense of Definition 24 corresponds to a deviation step of ρ^* in the sense of Definition 4[11]. Moreover the (H,i)-decomposition of ρ^* is finite (resp. infinite) iff σ'_i is finitely (resp. infinitely) deviating from $\bar\sigma_{\restriction h}$. Thus this play ρ^* belongs to $\mathbf{D}^{H,i}_\alpha(h'v')$, and by (12) we get $\lambda_i(h\cdot\langle\bar\sigma_{\restriction h}\rangle_v) > \lambda_i(hg\rho^*)$. Hence σ'_i is a profitable deviation for player i in $(\mathcal{G}_{\restriction h}, v)$, a contradiction with $\bar\sigma$ being an SPE.

Let α be a limit ordinal. By induction hypothesis $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}'_\beta(hv), \forall\beta < \alpha$. Therefore $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}'_\alpha(hv) = \bigcap_{\beta<\alpha}\mathbf{P}'_\beta(hv)$.

Let us now turn to the second statement of Theorem 27. Let $\rho \in \mathbf{P}'_{\beta_*}(v_0)$. By Proposition 14, it is enough to construct a very weak SPE $\bar\sigma$ with outcome ρ. The proof is very similar to the one of Lemma 21, where the construction of $\bar\sigma$ is done step by step thanks to a labeling γ of the histories. We briefly recall this proof and insist on the differences.

Initially, no history is labeled. We start with the play $\rho \in \mathbf{P}'_{\beta_*}(v_0)$, $\bar\sigma$ is partially defined such that $\langle\bar\sigma\rangle_{v_0} = \rho$, and $\gamma(h) = \rho$ for all non-empty prefixes h of ρ.

At the following steps, let $h'v'$ be a history that is not yet labeled, but such that h' has already been labeled. Suppose that $\mathrm{Last}(h') \in V_i$. By induction, $\gamma(h') = \langle\bar\sigma_{\restriction h}\rangle_v$ such that $hv \leq h'$, and $\langle\bar\sigma_{\restriction h}\rangle_v \in \mathbf{P}'_{\beta_*}(hv)$. Since $\mathbf{P}'_{\beta_*}(hv) = \mathbf{P}'_{\beta_*+1}(hv)$ by Proposition 26, we have $\langle\bar\sigma_{\restriction h}\rangle_v \notin \mathbf{E}'_{\beta_*}(hv)$. Therefore, with $H = \emptyset$ and $\mathbf{D}^{H,i}_{\beta_*}(h'v') = \mathbf{P}'_{\beta_*}(h'v')$, we know that there exists a play $\rho' \in \mathbf{P}'_{\beta_*}(h'v')$ such that $\lambda_i(h\cdot\langle\bar\sigma_{\restriction h}\rangle_v) \leq \lambda_i(h'\rho')$. Hence, we continue to construct $\bar\sigma$ such that $\langle\bar\sigma_{\restriction h'}\rangle_{v'} = \rho'$, and all non-empty prefixes g of ρ' are labeled by $\gamma(h'g) = \rho'$. And so on.

The constructed $\bar\sigma$ is a very weak SPE as in the proof of Lemma 21. ◀

The next proposition states that for cost functions that are upper-semicontinuous, sets $\mathbf{P}'_{\beta_*}(hv)$ and $\mathbf{P}_{\alpha_*}(hv)$ are all equal. This no longer the case as soon as one cost function is not upper-semicontinuous as shown by Examples 19 and 28.

▶ **Proposition 29.** *Let $(\mathcal{G}, v_0)$ be a quantitative game such that all its cost functions are upper-semicontinuous. Then for all $hv \in \mathrm{Hist}(v_0)$, $\mathbf{P}'_{\beta_*}(hv) = \mathbf{P}_{\alpha_*}(hv)$.*

Proof. Let us prove by induction on α that $\mathbf{P}'_\alpha(hv) \subseteq \mathbf{P}_\alpha(hv)$ for all hv. These two sets are equal for $\alpha = 0$. Let $\alpha+1$ be a successor ordinal and suppose that $\mathbf{P}'_\alpha(hv) \subseteq \mathbf{P}_\alpha(hv)$. We have $\mathbf{E}'_\alpha(hv) \supseteq \mathbf{E}_\alpha(hv)$, and thus $\mathbf{P}'_{\alpha+1}(hv) \subseteq \mathbf{P}_{\alpha+1}(hv)$, because $\mathbf{D}^{H,i}_\alpha(h'v') = \mathbf{P}'_\alpha(h'v')$ when $H = \emptyset$. For α being a limit ordinal, we easily have $\mathbf{P}'_\alpha(hv) \subseteq \mathbf{P}_\alpha(hv)$ by induction hypothesis.

Suppose now that $\mathbf{P}'_{\beta_*}(hv) \subsetneq \mathbf{P}_{\alpha_*}(hv)$ for some hv. Let $\rho \in \mathbf{P}_{\alpha_*}(hv) \setminus \mathbf{P}'_{\beta_*}(hv)$, and consider the initialized game $(\mathcal{G}', v'_0) = (\mathcal{G}_{\restriction h}, v)$. Notice that the sets $\mathbf{P}_\alpha(h'v')$ and $\mathbf{E}_\alpha(h'v')$ of this game $(\mathcal{G}', v'_0)$ are exactly the sets $\mathbf{P}_\alpha(hh'v')$ and $\mathbf{E}_\alpha(hh'v')$ of $(\mathcal{G}_{\restriction h}, v)$. By Theorem 18, there exists a weak SPE $\bar\sigma$ in $(\mathcal{G}', v'_0)$ with outcome ρ. Since the cost functions are upper-semicontinuous, σ is also an SPE by Proposition 14. Therefore, $\rho \in \mathbf{P}'_{\beta_*}(hv)$ by Theorem 27, which is a contradiction. ◀

[11] History $h'v'$ leads to one additional deviation step of ρ^* in the sense of Definition 4 (see (13)).

4 Quantitative Reachability Games

In this section, we focus on quantitative reachability games. Recall that in this case, the cost of a play for player i is the number of edges to reach his target set of vertices T_i (see Definition 2). Recall also that for quantitative reachability games, SPEs, weak SPEs, and very weak SPEs, are equivalent notions (see Corollary 15).

It is known that there always exists an SPE in quantitative reachability games [1, 6].

▶ **Theorem 30.** *Each quantitative reachability game $(\mathcal{G}, v_0)$ has an SPE.*

As SPEs and weak SPEs coincide in quantitative reachability games, we get the next result by Theorem 18.

▶ **Corollary 31.** *Let $(\mathcal{G}, v_0)$ be a quantitative reachability game. The sets $\mathbf{P}_{\alpha_*}(hv)$ are non-empty, for all $hv \in \mathrm{Hist}(v_0)$, and $\mathbf{P}_{\alpha_*}(v_0)$ is the set of outcomes of SPEs in $(\mathcal{G}, v_0)$.*

The proof provided for Theorem 30 is non constructive since it relies on topological arguments. Our main result is that one can algorithmically construct an SPE in $(\mathcal{G}, v_0)$ that is moreover finite-memory, thanks to the sets $\mathbf{P}_{\alpha_*}(hv)$.

▶ **Theorem 32.** *Each quantitative reachability initialized game $(\mathcal{G}, v_0)$ has a finite-memory SPE. Moreover there is an algorithm to construct such an SPE.*

We can also decide whether there exists a (finite-memory) SPE such that the cost of its outcome is component-wise bounded by a given constant vector.

▶ **Corollary 33.** *Let $(\mathcal{G}, v_0)$ be a quantitative reachability initialized game, and let $\bar{c} \in \mathbb{N}^{|\Pi|}$ be a given $|\Pi|$-uple of integers. Then one can decide whether there exists a (finite-memory) SPE $\bar{\sigma}$ such that $\lambda_i(\langle\bar{\sigma}\rangle_{v_0}) \leq c_i$ for all $i \in \Pi$.*

The main ingredients of the proof of Theorem 32 are the next ones; they will be detailed in the sequel of this section. We will give afterwards the proof of Corollary 33.

- Given α, the infinite number of sets $\mathbf{P}_\alpha(hv)$ can be replaced by the finite number of sets $\mathbf{P}^I_\alpha(v)$ where I is the set of players that did not reach their target set along history h.
- The fixpoint of Proposition 17 is reached with some natural number $\alpha_* \in \mathbb{N}$.
- Each $\mathbf{P}^I_\alpha(v)$ is a non-empty ω-regular set, thus containing a "lasso play" of the form $h \cdot g^\omega$.
- The lasso plays of each $\mathbf{P}^I_\alpha(v)$ allow to construct a finite-memory SPE.

The next lemma highlights a simple useful property of the cost functions λ_i used in quantitative reachability games. The proof is immediate.

▶ **Lemma 34.** *Let $i \in \Pi$ and $\rho \in \mathbf{P}_\alpha(hv)$. If player i did not reach his target set along history h, then $\lambda_i(h\rho) = \lambda_i(\rho) + |hv|$.*

The next proposition is a key result that will be used several times later on. It states that it is impossible to have plays in $\mathbf{P}_\alpha(hv)$ with arbitrarily large costs for player i, without having a play in $\mathbf{P}_\alpha(hv)$ with an infinite cost for player i.

▶ **Proposition 35.** *Consider $\mathbf{P}_\alpha(hv)$ and $i \in \Pi$. If for all $\rho \in \mathbf{P}_\alpha(hv)$, we have $\lambda_i(\rho) < +\infty$, then there exists c such that for all $\rho \in \mathbf{P}_\alpha(hv)$, we have $\lambda_i(\rho) \leq c$. The constant c only depends on $\mathbf{P}_\alpha(hv)$ and player i.*

Proof. Suppose that for all $n \in \mathbb{N}$, there exists $\rho_n \in \mathbf{P}_\alpha(hv)$ such that $\lambda_i(\rho_n) > n$. By König's lemma, there exists $\rho = \lim_{k\to\infty} \rho_{n_k}$ for some subsequence $(\rho_{n_k})_k$ of $(\rho_n)_n$. By definition of λ_i in quantitative reachability games, we get $\lambda_i(\rho) = +\infty$. Let us prove by induction on α that $\rho \in \mathbf{P}_\alpha(hv)$; this will establish Proposition 35.

Let $\alpha = 0$. As each $\rho_n \in \mathbf{P}_0(hv)$, then ρ_n is a play in $(\mathcal{G}_{\restriction h}, v)$ by definition of $\mathbf{P}_0(hv)$ (see (4) in Definition 16). Therefore ρ is also a play in $(\mathcal{G}_{\restriction h}, v)$, and $\rho \in \mathbf{P}_0(hv)$.

Let $\alpha + 1$ be a successor ordinal. As for all n, $\rho_n \in \mathbf{P}_{\alpha+1}(hv) \subseteq \mathbf{P}_\alpha(hv)$, then $\rho \in \mathbf{P}_\alpha(hv)$ by induction hypothesis. Let us prove that $\rho \in \mathbf{P}_{\alpha+1}(hv)$. Suppose on the contrary that $\rho \in \mathbf{E}_\alpha(hv)$ (see Definition 16 and Figure 5). Then there exists a history $h' \in \mathrm{Hist}_j$ with $j \in \Pi$ and $hv \leq h' < h\rho$, there exists a vertex v' with $h'v' \not< h\rho$, such that $\forall \rho' \in \mathbf{P}_\alpha(h'v')$:

$$\lambda_j(h\rho) > \lambda_j(h'\rho'). \tag{14}$$

It follows that player j did not reach his target set along h'. Hence by Lemma 34, we have $\lambda_j(h\rho) = \lambda_j(\rho) + |hv|$ and $\lambda_j(h'\rho') = \lambda_j(\rho') + |h'v'|$. By (14), $\lambda_j(\rho')$ is bounded. Hence by induction hypothesis with $\mathbf{P}_\alpha(h'v')$ and $j \in \Pi$, there exists a constant c such that $\lambda_j(\rho') \leq c$, $\forall \rho' \in \mathbf{P}_\alpha(h'v')$.

Suppose first that $\lambda_j(\rho) < +\infty$. Then, since $\rho = \lim_{k\to\infty} \rho_{n_k}$, it follows that for a large enough n_k, the plays ρ and ρ_{n_k} share a long common prefix on which player j reaches its target set, i.e. $\lambda_j(\rho) = \lambda_j(\rho_{n_k})$. It follows that with the same history $h'v'$ as above, by (14) and Lemma 34, we have $\lambda_j(h\rho_{n_k}) > \lambda_j(h'\rho')$, $\forall \rho' \in \mathbf{P}_\alpha(h'v')$, showing that $\rho_{n_k} \in \mathbf{E}_\alpha(hv)$, a contradiction.

Suppose next that $\lambda_j(\rho) = +\infty$. Then, given $c' = |h'v'| - |hv| + c$, we can choose a large enough n_k such that the plays ρ and ρ_{n_k} share a common prefix of length at least c'. Moreover, as $\lambda_j(\rho) = +\infty$, player j does not reach its target set along this prefix, i.e. $\lambda_j(\rho_{n_k}) > c'$. Therefore, using the same history $h'v'$ as above, by (14) and Lemma 34, we have $\lambda_j(h\rho_{n_k}) > |hv| + c' = |h'v'| + c \geq \lambda_j(h'\rho')$, $\forall \rho' \in \mathbf{P}_\alpha(h'v')$. This shows that $\rho_{n_k} \in \mathbf{E}_\alpha(hv)$, again a contradiction.

Let α be a limit ordinal. As for all n, $\rho_n \in \mathbf{P}_\alpha(hv) = \bigcap_{\beta<\alpha} \mathbf{P}_\beta(hv)$ (see (6) in Definition 16), then $\rho \in \mathbf{P}_\beta(hv), \forall \beta < \alpha$, by induction hypothesis. Hence $\rho \in \mathbf{P}_\alpha(hv) = \bigcap_{\beta<\alpha} \mathbf{P}_\beta(hv)$.

We have just shown that if all plays $\rho \in \mathbf{P}_\alpha(hv)$ have a cost $\lambda_i(\rho) < +\infty$, then there exists c such that $\lambda_i(\rho) \leq c$ for all such ρ. This constant c depends on $\mathbf{P}_\alpha(hv)$ and player i. ◀

As a consequence of Proposition 35, we have that $\sup\{\lambda_i(\rho) \mid \rho \in \mathbf{P}_\alpha(hv)\}$ is equal to $\max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}_\alpha(hv)\}$, and that this maximum belongs to $\mathbb{N} \cup \{+\infty\}$.

4.1 Sets $\mathbf{P}^I_\alpha(v)$

Let $(\mathcal{G}, v_0)$ be a quantitative reachability game. Given a history $h = h_0 \dots h_n$ in $(\mathcal{G}, v_0)$, we denote by $I(h)$ the set of players i such that $\forall k$, $0 \leq k \leq n$, we have $h_k \notin T_i$. In other words $I(h)$ is the set of players that did not reach their target set along history h. If h is empty, then $I(h) = \Pi$. The next lemma indicates that sets $\mathbf{P}_\alpha(hv)$ only depend on v and $I(h)$, and thus not on h (we do no longer take care of players that have reached their target set along h).

▶ **Lemma 36.** *For $h_1v, h_2v \in \mathrm{Hist}(v_0)$, if $I(h_1) = I(h_2)$, then $\mathbf{P}_\alpha(h_1v) = \mathbf{P}_\alpha(h_2v)$ for all α.*

Proof. The proof is by induction on α. By definition, we have $P_0(h_1v) = P_0(h_2v)$.

Suppose that $\alpha + 1$ is a successor ordinal. By induction hypothesis, $\mathbf{P}_\alpha(h_1v) = \mathbf{P}_\alpha(h_2v)$. Let us prove that $\mathbf{E}_\alpha(h_1v) = \mathbf{E}_\alpha(h_2v)$ which will imply that $\mathbf{P}_{\alpha+1}(h_1v) = \mathbf{P}_{\alpha+1}(h_2v)$. If

$\rho \in \mathbf{E}_\alpha(h_1 v)$, it means that there exists a history $h'_1 = h_1 g \in \mathrm{Hist}_i$ with $h'_1 < h_1\rho$, there exists a vertex v' with $h'_1 v' \not< h_1\rho$, such that $\forall \rho' \in \mathbf{P}_\alpha(h'_1 v')$, we have $\lambda_i(h_1\rho) > \lambda_i(h'_1\rho')$, i.e.

$$\lambda_i(\rho) > \lambda_i(g\rho'). \tag{15}$$

In particular, $i \in I(h'_1)$. Let us consider the history $h'_2 = h_2 g$. By hypothesis, $I(h_1) = I(h_2)$, and therefore $I(h'_1) = I(h'_2)$ and $i \in I(h'_2)$. Thus by induction hypothesis $\mathbf{P}_\alpha(h'_1 v') = \mathbf{P}_\alpha(h'_2 v')$. It follows that for $\forall \rho' \in \mathbf{P}_\alpha(h'_2 v')$, we have $\lambda_i(h_2\rho) > \lambda_i(h_2 g\rho') = \lambda_i(h'_2\rho')$ by (15), and then $\rho \in \mathbf{E}_\alpha(h_2 v)$. Symmetrically, if $\rho \in \mathbf{E}_\alpha(h_2 v)$, then $\rho \in \mathbf{E}_\alpha(h_1 v)$. We can conclude that $\mathbf{P}_{\alpha+1}(h_1 v) = \mathbf{P}_{\alpha+1}(h_2 v)$.

Suppose that α is a limit ordinal. As $\mathbf{P}_\alpha(h_1 v) = \bigcap_{\beta<\alpha} \mathbf{P}_\beta(h_1 v)$, and $\mathbf{P}_\beta(h_1 v) = \mathbf{P}_\beta(h_2 v)$ by induction hypothesis, it follows that $\mathbf{P}_\alpha(h_1 v) = \mathbf{P}_\alpha(h_2 v)$. ◀

Thanks to this lemma, we can introduce the next definitions.

▶ **Definition 37.** Let $(\mathcal{G}, v_0)$ be a quantitative reachability initialized game. Let $I \subseteq \Pi$ be such that $I = I(h)$ for some $h \in \mathrm{Hist}(v_0)$. We denote by

- $\mathbf{P}^I_\alpha(v)$ the set $\mathbf{P}_\alpha(hv)$, and by
- $\mathbf{E}^I_\alpha(v)$ the set $\mathbf{E}_\alpha(hv)$.

In particular, $\mathbf{P}^\Pi_\alpha(v_0) = \mathbf{P}_\alpha(v_0)$ and $\mathbf{E}^\Pi_\alpha(v_0) = \mathbf{E}_\alpha(v_0)$.

Given α, the infinite number of sets $\mathbf{P}_\alpha(hv)$ can thus be replaced by the finite number of sets $\mathbf{P}^I_\alpha(v)$. Moreover, Proposition 35 can be rephrased as follows.

▶ **Corollary 38.** *Consider $\mathbf{P}^I_\alpha(v)$ and $i \in I$. If for all $\rho \in \mathbf{P}^I_\alpha(v)$, we have $\lambda_i(\rho) < +\infty$, then there exists c such that for all $\rho \in \mathbf{P}^I_\alpha(v)$, we have $\lambda_i(\rho) \leq c$. The constant c only depends on α, I, v, and i.*

Proof. Let h be such that $I = I(h)$. Consider $\mathbf{P}_\alpha(hv) = \mathbf{P}^I_\alpha(v)$, and $i \in I$. By Proposition 35, if for all $\rho \in \mathbf{P}_\alpha(hv)$, $\lambda_i(\rho) < +\infty$, then there exists c (depending on $\mathbf{P}_\alpha(hv)$ and i) such that for all $\rho \in \mathbf{P}_\alpha(hv)$, $\lambda_i(\rho) \leq c$. By Lemma 36, c depends on α, I, v, and i. ◀

As a consequence of Corollary 38, we give the next definition that indicates the maximum costs for plays in $\mathbf{P}^I_\alpha(v)$.

▶ **Definition 39.** Given $\mathbf{P}^I_\alpha(v)$, we define $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$ such that

$$\Lambda_i(\mathbf{P}^I_\alpha(v)) = \begin{cases} -1 & \text{if } i \notin I, \\ \max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v)\} & \text{if } i \in I. \end{cases}$$

In this definition, -1 indicates that player i has already visited his target set T_i, and the max belongs to $\mathbb{N} \cup \{+\infty\}$.

4.2 Fixpoint with $\alpha_* \in \mathbb{N}$

In this section, we aim at proving that the fixpoint (when computing the sets $\mathbf{P}^I_\alpha(v)$, see Proposition 17) is reached in a finite number of steps, that is $\alpha_* \in \mathbb{N}$.

We first need to introduce some notions about the sets $\mathbf{P}^I_\alpha(v)$. Let $\rho = \rho_0\rho_1\ldots \in \mathbf{P}^I_\alpha(v)$. We use a map χ that *decorates* each ρ_n by some set $J \subseteq \Pi$. The aim of the decoration $\chi(\rho_n)$ is to indicate at vertex ρ_n, which players of I did not reach their target set along $\rho_{<n}$. More precisely, $\chi(\rho_n) = I \cap I(\rho_{<n})$. In particular $\chi(\rho_0) = I \cap \Pi = I$.

Let $\mathbf{P}^I_\alpha(v)$ and $(v,v') \in E$. We now adapt Definition 39 to mention the maximum costs for plays in $\mathbf{P}^I_\alpha(v)$ *starting with edge* (v,v'). We define $\overline{\Lambda}(\mathbf{P}^I_\alpha(v),v')$ as follows:

$$\Lambda_i(\mathbf{P}^I_\alpha(v),v') = \begin{cases} -1 & \text{if } i \notin I, \\ \max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v) \text{ and } \rho_0\rho_1 = vv'\} & \text{if } i \in I. \end{cases} \tag{16}$$

In this definition, the max is equal to -1 when the set $\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v) \text{ and } \rho_0\rho_1 = vv'\}$ is empty.[12]

The sequence $(\overline{\Lambda}(\mathbf{P}^I_\alpha(v),v'))_\alpha$ is nonincreasing for the usual component-wise ordering over $(\mathbb{N} \cup \{-1,+\infty\})^\Pi$ since $\mathbf{P}^I_\alpha(v)$ is nonincreasing for the inclusion by definition. Therefore it reaches a fixpoint that we want to relate to the fixpoint $\mathbf{P}^I_{\alpha_*}(v)$ of Proposition 17. This is done in the following lemma.

▶ **Lemma 40.**
- *If* $\mathbf{P}^I_\alpha(v) = \mathbf{P}^I_{\alpha+1}(v)$*, then for all* $(v,v') \in E$, $\overline{\Lambda}(\mathbf{P}^I_\alpha(v),v') = \overline{\Lambda}(\mathbf{P}^I_{\alpha+1}(v),v')$.
- *If* $\mathbf{P}^I_\alpha(v) \neq \mathbf{P}^I_{\alpha+1}(v)$*, then there exist* $J \subseteq \Pi$ *and* $(u,u') \in E$ *such that* $\overline{\Lambda}(\mathbf{P}^J_\alpha(u),u') \neq \overline{\Lambda}(\mathbf{P}^J_{\alpha+1}(u),u')$.

Proof. The first statement is immediate from definition of $\overline{\Lambda}$. Let us prove the second statement. Consider $\rho = \rho_0\rho_1\ldots \in \mathbf{E}^I_\alpha(v)$. Then there exist $i \in \Pi$, $n \in \mathbb{N}$ and $v' \neq \rho_{n+1}$ with $\rho_n \in V_i$, $\chi(\rho_n) = J$, $\chi(\rho_{n+1}) = J'$, such that $\forall \rho' \in \mathbf{P}^{J'}_\alpha(v')$ we have $\lambda_i(\rho) > \lambda_i(\rho_0\ldots\rho_n\rho')$ or equivalently (by Lemma 34)

$$\lambda_i(\rho) - (n+1) > \lambda_i(\rho'). \tag{17}$$

(see the definition of $\mathbf{E}_\alpha(hv)$ with $I(h) = I$ in Definition 16 and Figure 9). Notice that $i \in J'$. Let us prove that $\overline{\Lambda}(\mathbf{P}^J_\alpha(u),u') \neq \overline{\Lambda}(\mathbf{P}^J_{\alpha+1}(u),u')$ with $u = \rho_n$ and $u' = \rho_{n+1}$. As $\rho \in \mathbf{P}^I_\alpha(v)$,

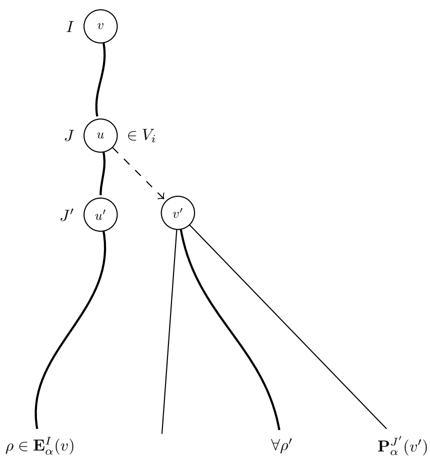

Figure 9 $\rho \in \mathbf{E}^I_\alpha(v)$, with $\rho_n = u$ and $\rho_{n+1} = u'$.

then $\rho_{\geq n} \in \mathbf{P}^J_\alpha(u)$ by Lemma 22. As $\lambda_i(\rho_{\geq n}) = \lambda_i(\rho) - n$, this implies that

$$\Lambda_i(\mathbf{P}^J_\alpha(u),u') \geq \lambda_i(\rho) - n. \tag{18}$$

Let $\varrho \in \mathbf{P}^J_\alpha(u)$ be such that ϱ starts with edge (u,u') and has maximal cost $\Lambda_i(\mathbf{P}^J_\alpha(u),u')$. One gets

$$\lambda_i(\varrho) = \Lambda_i(\mathbf{P}^J_\alpha(u),u') \geq \lambda_i(\rho) - n > \lambda_i(\rho_n\rho')$$

[12] Notice that as $\mathbf{P}^I_\alpha(v)$ is non-empty, there exists some $(v,v') \in E$ such that this set is non-empty.

by (17) and (18). By considering the set $\mathbf{P}^{J'}_\alpha(v')$ in Figure 9, it follows that $\varrho \in \mathbf{E}^J_\alpha(u)$ for all such plays ϱ. Hence $\mathbf{P}^J_{\alpha+1}(u) \subsetneq \mathbf{P}^J_\alpha(u)$ and $\Lambda_i(\mathbf{P}^J_{\alpha+1}(u), u') < \Lambda_i(\mathbf{P}^J_\alpha(u), u')$. This completes the proof. ◀

We are now able to prove that the ordinal α_* of Proposition 17 is an integer.

▶ **Corollary 41.** *There exists an integer α_* such that $\mathbf{P}^I_{\alpha_*}(v) = \mathbf{P}^I_{\alpha_*+1}(v)$ for all $v \in V$ and $I \subseteq \Pi$.*

Proof. Notice that there is a finite number of sequences $(\overline{\Lambda}(\mathbf{P}^I_\alpha(v), v'))_\alpha$ since they depend on $I \subseteq \Pi$ and $(v, v') \in E$. As the component-wise ordering over $(\mathbb{N} \cup \{-1, +\infty\})^\Pi$ is a well-quasi-ordering and all these sequences are nonincreasing, there exists an integer (and not only an ordinal) α'_* such that $\overline{\Lambda}(\mathbf{P}^I_{\alpha'_*}(v), v') = \overline{\Lambda}(\mathbf{P}^I_{\alpha'_*+1}(v), v')$ for all $I \subseteq \Pi$ and $(v, v') \in E$. By Lemma 40, we get that $\alpha_* \le \alpha'_*$, showing that $\alpha_* \in \mathbb{N}$. ◀

4.3 The sets $\mathbf{P}^I_\alpha(v)$ are ω-regular

In this section, we prove that each set $\mathbf{P}^I_\alpha(v)$ is ω-regular. Instead of providing the construction of a Büchi automaton (which would lead to many technical details), we prefer to show that each set $\mathbf{P}^I_\alpha(v)$ is MSO-definable. It is well-known that a set of ω-words is ω-regular iff it is MSO-definable, by Büchi theorem [17]. Moreover from the Büchi automaton, one can construct an equivalent MSO-sentence, and conversely. One can also decide whether an MSO-sentence is satisfiable [17]. We recall that MSO-logic uses:

- variables $x, y, \dots$ ($X, Y, \dots$ resp.) to describe a position (a set of positions resp.) in an ω-word ρ, and relations $X(x)$ to mention that x belongs to X,
- relations $Q_u(x)$, $u \in V$, to mention that such vertex u is at position x of ρ,
- relations $x < y$ and $Succ(x, y)$ to mention that position y is after position x, and position y is successor of position x respectively,
- connectives $\vee, \wedge, \neg$ and quantifiers $\exists x$, $\forall x$, $\exists X$, $\forall X$

Recall that constants $0, 1, \dots$ are definable. We will use notation $x + 1$ (and more generally $x + c$, with c a constant) instead of $Succ(x, y)$.

▶ **Proposition 42.** *Each $\mathbf{P}^I_\alpha(v)$ is an ω-regular set.*

We begin with a lemma that states that if $\mathbf{P}^I_\alpha(v)$ is ω-regular, then the maximum of its costs is computable.

▶ **Lemma 43.** *If $\mathbf{P}^I_\alpha(v)$ is MSO-definable, then $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$ is computable.*

Proof. Before proving this lemma, we need to establish two properties. The first one states that one can decide whether $\mathbf{P}^I_\alpha(v)$ has a play ρ with a given cost for a given player. The second one states that when $\Lambda_i(\mathbf{P}^I_\alpha(v))$ is finite, then this number is bounded by the number of states of a Büchi automaton accepting $\mathbf{P}^I_\alpha(v)$.

(i) Let $c \in \mathbb{N} \cup \{+\infty\}$ and $i \in I$. Let ϕ be an MSO-sentence defining $\mathbf{P}^I_\alpha(v)$. Let us show that one can decide whether $\mathbf{P}^I_\alpha(v)$ has a play ρ with cost $\lambda_i(\rho) = c$. There exists an MSO-sentence φ expressing that $\lambda_i(\rho) = c$. Indeed, if $c = +\infty$, then φ is the sentence $\forall x \cdot \neg(\vee_{u \in T_i} Q_u(x))$, and if $c < +\infty$, it is the sentence $(\forall x < c \cdot \neg(\vee_{u \in T_i} Q_u(x))) \wedge (\vee_{u \in T_i} Q_u(c))$. Therefore one can decide whether the MSO-sentence $\phi \wedge \varphi$ is satisfiable by some play ρ.

(ii) Let $i \in I$ and suppose that $\Lambda_i(\mathbf{P}^I_\alpha(v)) < +\infty$. Let $\mathcal{B}$ be a Büchi automaton accepting $\mathbf{P}^I_\alpha(v)$. We now show that $\Lambda_i(\mathbf{P}^I_\alpha(v)) < n$ where n is the number of states of $\mathcal{B}$. Assume the contrary and consider an accepting run $r = r_0 r_1 \dots$ of $\mathcal{B}$ on a play $\rho = \rho_0 \rho_1 \dots \in \mathbf{P}^I_\alpha(v)$ with $\lambda_i(\rho) = \Lambda_i(\mathbf{P}^I_\alpha(v)) \ge n$. The prefix $r_{\le n}$ of r has a cycle $r_k \dots r_l$ with $0 \le k < l \le n$

and $r_k = r_l$. This cycle can be repeated once, while keeping an accepting run labeled by $\rho' = \rho_0 \dots (\rho_k \dots \rho_{l-1})^2 \rho_{\geq l}$. As $\lambda_i(\rho) \geq n$, it follows that $\lambda_i(\rho') = \Lambda_i(\mathbf{P}^I_\alpha(v)) + (l-k)$. Therefore we get a contradiction with $\lambda_i(\rho) = \Lambda_i(\mathbf{P}^I_\alpha(v))$.

Let us prove the lemma. By definition $\Lambda_i(\mathbf{P}^I_\alpha(v))$ equals -1 if $i \notin I$, and is thus computable in this case. Let $i \in I$. By *(i)*, one can decide whether $\Lambda_i(\mathbf{P}^I_\alpha(v)) = +\infty$. In case of a positive answer, $\Lambda_i(\mathbf{P}^I_\alpha(v))$ is thus computable. If the answer is negative, as $\Lambda_i(\mathbf{P}^I_\alpha(v)) < n$ by *(ii)*, we can similarly test whether $\Lambda_i(\mathbf{P}^I_\alpha(v)) = c$ by considering decreasing constants c from $n-1$ to 0. This prove that $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$ is computable. ◀

Proof of Proposition 42. Let us prove that each set $\mathbf{P}^I_\alpha(v)$ is MSO-definable by induction on α.

For $\alpha = 0$, recall that $\mathbf{P}^I_0(v)$ is the set of plays starting with v. The required sentence is thus $Q_v(0) \wedge \forall x \cdot \vee_{(u,u') \in E} (Q_u(x) \wedge Q_{u'}(x+1))$.

Let $\alpha \in \mathbb{N}$ be a fixed integer. By induction hypothesis, each set $\mathbf{P}^I_\alpha(v)$ is MSO-definable, and by Lemma 43, $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$ is computable. These sets and constants can be considered as fixed. Let us prove that $\mathbf{E}^I_\alpha(v)$ is MSO-definable. It will follow that $\mathbf{P}^I_{\alpha+1}(v)$ is also MSO-definable. Thanks to $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$, the definition of $\rho \in \mathbf{E}^I_\alpha(v)$ can be rephrased as follows: there exist $n \in \mathbb{N}$, $i \in I$, and $u, u', v' \in V$ with $u' \neq v'$, $(u, v') \in E$, such that $\rho_n = u \in V_i$, $\rho_{n+1} = u'$, $\chi(\rho_{n+1}) = J'$, and

$$\lambda_i(\rho) > \Lambda_i(\mathbf{P}^{J'}_\alpha(v')) + (n+1) \tag{19}$$

(see Figure 9). Notice that (19) implies that $i \in J'$ and $\Lambda_i(\mathbf{P}^{J'}_\alpha(v')) < +\infty$. Moreover $\Lambda_i(\mathbf{P}^{J'}_\alpha(v'))$ is a fixed integer.

Let us provide an MSO-sentence ψ defining $\mathbf{E}^I_\alpha(v)$. The next sentence $\phi_{J',n}$ expresses that $J' \subseteq I$ is the subset of players of I that did not visit their target set along $\rho_{\leq n}$:

$$\phi_{J',n} = \big(\forall x \cdot (x \leq n) \rightarrow \neg(\vee_{j \in J'} \vee_{r \in T_j} Q_r(x))\big) \wedge \big(\wedge_{j \in I \setminus J'} \exists x \leq n \cdot \vee_{r \in T_j} Q_r(x)\big)\,.$$

The next sentence $\varphi_{J',n,v',i}$ expresses that if player i visits its target set along ρ, it is after $\Lambda_i(\mathbf{P}^{J'}_\alpha(v')) + n + 1$ edges from ρ_0:

$$\varphi_{J',n,v',i} = \forall x \cdot \Big(\vee_{r \in T_i} Q_r(x) \rightarrow \Big(x > \Lambda_i(\mathbf{P}^{J'}_\alpha(v')) + n + 1\Big)\Big)\,.$$

Notice that in the previous formula, $\Lambda_i(\mathbf{P}^{J'}_\alpha(v'))$ is a constant since $\mathbf{P}^{J'}_\alpha(v')$ is a fixed set. The required formula ψ is then the following one:

$$\exists n \cdot \bigvee_{\substack{u,u' \neq v' \in V \\ (u,v') \in E}} \; \bigvee_{J' \subseteq I} \; \bigvee_{\substack{i \in J', u \in V_i \\ \Lambda_i(\mathbf{P}^{J'}_\alpha(v')) < +\infty}} (Q_u(n) \wedge Q_{u'}(n+1) \wedge \phi_{J',n} \wedge \varphi_{J',n,v',i})\,.$$

◀

By Proposition 42, the next corollary states that one can effectively extract a lasso play from $\mathbf{P}^I_\alpha(v)$ that has a maximal cost for a given player i.

▶ **Corollary 44.** *For all $i \in I$, each set $\mathbf{P}^I_{\alpha_*}(v)$ has a computable lasso play $h \cdot g^\omega$ with $\lambda_i(h \cdot g^\omega) = \Lambda_i(\mathbf{P}^I_{\alpha_*}(v))$. This play depends on i, I, and v.*

Proof. Part *(i)* of the proof of Lemma 43 indicates that the set of plays $\rho \in \mathbf{P}^I_{\alpha_*}(v)$ with maximal cost $\lambda_i(\rho) = \Lambda_i(\mathbf{P}^I_{\alpha_*}(v))$ is ω-regular. Therefore, from a Büchi automaton accepting this set, we can extract an accepted lasso play of the form $h \cdot g^\omega$ with the required cost. Such a play depends on i, I, and v (α_* is fixed). ◀

4.4 Construction of a Finite-Memory SPE

Thanks to the results of Sections 4.1-4.3, we have all the ingredients to prove that each quantitative reachability game has a computable finite-memory SPE.

Proof of Theorem 32. Let $(\mathcal{G}, v_0)$ be a quantitative reachability game. Let us summarize the results obtained previously. By Corollary 31, each set $\mathbf{P}^I_{\alpha_*}(v)$ is non-empty with $v \in V$ and $I = I(h)$ for some history $hv \in \text{Hist}(v_0)$, and $\mathbf{P}^\Pi_{\alpha_*}(v_0)$ contains all the outcomes of SPEs in $(\mathcal{G}, v_0)$. By Corollary 41 and Proposition 42, we know that $\alpha_* \in \mathbb{N}$ and each $\mathbf{P}^I_{\alpha_*}(v)$ is an ω-regular set that can be constructed. Finally by Corollary 44, for all $i \in I$, one can construct a lasso play $h_{i,I,v}\cdot(g_{i,I,v})^\omega \in \mathbf{P}^I_{\alpha_*}(v)$ with maximal cost $\lambda_i(h_{i,I,v}\cdot(g_{i,I,v})^\omega) = \Lambda_i(\mathbf{P}^I_{\alpha_*}(v))$.

We now show how to construct a finite-memory SPE $\bar{\sigma}$ from the finite set of lasso plays $h_{i,I,v}\cdot(g_{i,I,v})^\omega$. The procedure is similar to the one developed in the proof of Theorem 18 and more particularly of Lemma 21. We indicate how to adapt the proof of this lemma. Again the construction of $\bar{\sigma}$ is done step by step, thanks to a labeling γ of the non-empty histories.

Initially, none of the histories is labeled. We start with history v_0 and with any play $h_{i,\Pi,v_0}\cdot(g_{i,\Pi,v_0})^\omega \in \mathbf{P}^\Pi_{\alpha_*}(v_0)$, $i \in \Pi$. The strategy profile $\bar{\sigma}$ is partially defined such that $\langle\bar{\sigma}\rangle_{v_0} = h_{i,\Pi,v_0}\cdot(g_{i,\Pi,v_0})^\omega$, and the non-empty prefixes h of $h_{i,\Pi,v_0}\cdot(g_{i,\Pi,v_0})^\omega$ are all labeled with $\gamma(h) = (i, \Pi, v_0)$.

At the following steps, we consider a history $h'v'$ that is not yet labeled, but such that h' has already been labeled. By induction, $\gamma(h') = (j, I, v)$ and there exists $hv \leq h'$ such that $\langle\bar{\sigma}_{\restriction h}\rangle_v = h_{j,I,v}\cdot(g_{j,I,v})^\omega$. Suppose that $\text{Last}(h') \in V_i$ and $I(h') = J'$, the proof of Lemma 21 requires to choose[13] a play $\rho' \in \mathbf{P}^{J'}_{\alpha_*}(v')$ such that $\lambda_i(h'\rho') \geq \lambda_i(h\langle\bar{\sigma}_{\restriction h}\rangle_v)$ (see (8)). We simply choose $\rho' = h_{i,J',v'}\cdot(g_{i,J',v'})^\omega$ that has maximal cost $\lambda_i(h_{i,J',v'}\cdot(g_{i,J',v'})^\omega) = \Lambda_i(\mathbf{P}^{J'}_{\alpha_*}(v'))$. Then we continue the construction of $\bar{\sigma}$ such that $\langle\bar{\sigma}_{\restriction h'}\rangle_{v'} = h_{i,J',v'}\cdot(g_{i,J',v'})^\omega$, and for all non-empty prefixes g of $h_{i,J',v'}\cdot(g_{i,J',v'})^\omega$, we define $\gamma(h'g) = (i, J', v')$.

By the proof of Lemma 21, the strategy profile $\bar{\sigma}$ is an SPE. It is finite-memory since for all $h \in \text{Hist}_i$, $\sigma_i(h)$ only depends on $\gamma(h) = (j, I, v)$ and $h_{j,I,v}\cdot(g_{j,I,v})^\omega$. There is a finite number of lasso plays $h_{j,I,v}\cdot(g_{j,I,v})^\omega$, and $\gamma(h)$ (as well as $I(h)$) can be computed inductively as follows. Initially, $I(v_0) = \Pi$, and $\gamma(v_0) = (i, \Pi, v_0)$ for some chosen $i \in \Pi$. Let $h' \in \text{Hist}_i$ and suppose that $I(h') = J'$ and $\gamma(h') = (j, I, v)$. Then $I(h'v') = J' \setminus \{i \mid v' \in T_i\}$. If $h'v'$ respects $h_{j,I,v}(g_{j,I,v})^\omega$, i.e. $\sigma_i(h') = v'$, then $\gamma(h'v') = (j, I, v)$. Otherwise $\gamma(h'v') = (i, J', v')$ with (i, J', v') computed as in the previous paragraph. ◀

4.5 Constrained Existence

It remains to prove the decidability of the constrained existence of SPE for quantitative reachability games, as announced in Corollary 33. This result is easily proved on the basis of some previous properties.

Proof of Corollary 33. Let $(\mathcal{G}, v_0)$ be a game and let $\bar{c} \in \mathbb{N}^{|\Pi|}$ be a constant vector. In the proof of Lemma 43, we have seen that there exists an MSO-sentence expressing that play ρ has a fixed cost $\lambda_i(\rho) = c_i$. Similarly, one can express that $\lambda_i(\rho) \leq c_i$ by the next sentence φ_i: $\exists x \leq c_i \cdot (\vee_{u \in T_i} Q_u(x))$. By Proposition 42, we know that the set $\mathbf{P}^\Pi_{\alpha_*}(v_0)$ of outcomes of SPEs in $(\mathcal{G}, v_0)$ is an ω-regular set, and that one can construct an MSO-sentence ϕ defining it. Therefore the set of outcomes of SPEs with a cost component-wise bounded by $\bar{c}$ is definable by $\wedge_{i\in\Pi}\varphi_i \wedge \phi$, and is then ω-regular. Moreover, one can decide whether this set is non-empty.

---

[13] This proof states that such a play always exists.

In case of positive answer, it contains a lasso play $h \cdot g^\omega$. Exactly as done in Section 4.4, one can construct a finite-memory SPE $\bar{\sigma}$ such that $\langle \bar{\sigma} \rangle_{v_0} = h \cdot g^\omega$. This concludes the proof. ◀

5 Games with Prefix-independent Regular Cost Functions

In this section, we present a class of games for which it is decidable whether there exists a weak SPE.[14] The hypotheses are general conditions on the cost functions λ_i, $i \in \Pi$: each function λ_i must be prefix-independent (see Definition 3), λ_i has to use a finite number of values (gathered in set C_i), and the set of plays ρ with a given cost $\lambda_i(\rho) = c_i$ must be ω-regular.

▶ **Theorem 45.** *Let $(\mathcal{G}, v_0)$ be an initialized game such that:*

- *each cost function λ_i is prefix-independent, and with finite range $C_i \subset \mathbb{Q}$,*
- *for all $i \in \Pi$, $c_i \in C_i$, and $v \in V$, the set of plays ρ in $(\mathcal{G}, v)$ with $\lambda_i(\rho) = c_i$ is an ω-regular set.*

Then one can decide whether $(\mathcal{G}, v_0)$ has a weak SPE $\bar{\sigma}$ (resp. such that $\lambda_i(\langle \bar{\sigma} \rangle_{v_0}) \leq c_i$ forall i for given $c_i \in C_i$, $i \in \Pi$). In case of positive answer, one can construct such a finite-memory weak SPE.

For example, the hypotheses of this theorem are satisfied by the liminf games and the limsup games; they are also satisfied by the game of Example 12. We will see that the proof of this decidability result shares similar points with the proof given in the previous section for quantitative reachability games. Again, we will use the Folk Theorem for weak SPEs (see Theorem 18) to prove this result. The main steps of the proof are the following ones.

- Given α, the infinite number of sets $\mathbf{P}_\alpha(hv)$ can be replaced by the finite number of sets $\mathbf{P}_\alpha(v)$.
- The fixpoint of Proposition 17 is reached with some natural number $\alpha_* \in \mathbb{N}$.
- Each $\mathbf{P}_\alpha(v)$ is an ω-regular set. Therefore there exists an algorithm to construct the sets $\mathbf{P}_{\alpha_*}(v)$ for all $v \in V$, and thus to decide whether they are all non-empty. For given constants $c_i \in C_i$, $i \in \Pi$, one can also decide whether $\mathbf{P}_{\alpha_*}(v_0)$ has a play ρ with bounded cost $\lambda_i(\rho) \leq c_i$ for all i.
- In case of positive answer, some lasso plays of the sets $\mathbf{P}_{\alpha_*}(v)$ allow to construct a finite-memory weak SPE (resp. with bounded cost).

To establish Theorem 45, we prove a series of lemmas. The first lemma states that $\mathbf{P}_\alpha(hv)$ is independent of h. There is thus a finite number of sets $\mathbf{P}_\alpha(v)$, $v \in V$, to study.

▶ **Lemma 46.** *$\mathbf{P}_\alpha(hv) = \mathbf{P}_\alpha(v)$ for all $hv \in \mathrm{Hist}(v_0)$.*

Proof. The proof can be easily done by induction on α. It uses the definition of $\mathbf{E}_\alpha(hv)$ and the hypothesis of Theorem 45 that each cost function λ_i is prefix-independent. ◀

As each cost function λ_i is supposed to have finite range in Theorem 45, we can give the next definition that indicates the maximum costs for plays in $\mathbf{P}_\alpha(v)$ (resp. starting with vv', for some given $(v, v') \in E$). Recall that a similar definition was given in case of quantitative reachability games (see Definition 39 and (16)).

▶ **Definition 47.** Given $\mathbf{P}_\alpha(v)$ and $(v, v') \in E$, we define for each $i \in \Pi$:

[14] Contrarily to quantitative reachability games, we do not know if a weak SPE always exists for games in this class.

- $\Lambda_i(\mathbf{P}_\alpha(v)) = \max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}_\alpha(v)\}$.
- $\Lambda_i(\mathbf{P}_\alpha(v), v') = \max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}_\alpha(v) \text{ and } \rho_0\rho_1 = vv'\}$.

In this definition, the max is equal to $-\infty$ if it applies to an empty set.

The sequence $(\overline{\Lambda}(\mathbf{P}_\alpha(v), v'))_\alpha$ is nonincreasing for the component-wise ordering over $(\mathbb{Q} \cup \{-\infty\})^{|\Pi|}$.[15] Therefore it reaches a fixpoint as for $(\mathbf{P}_\alpha(v))_\alpha$. The following lemma relates these sequences.

▶ **Lemma 48.**
- *If $\mathbf{P}_\alpha(v) = \mathbf{P}_{\alpha+1}(v)$, then for all $(v, v') \in E$, $\overline{\Lambda}(\mathbf{P}_\alpha(v), v') = \overline{\Lambda}(\mathbf{P}_{\alpha+1}(v), v')$.*
- *If $\mathbf{P}_\alpha(v) \neq \mathbf{P}_{\alpha+1}(v)$, then there exists $(u, u') \in E$ such that $\overline{\Lambda}(\mathbf{P}_\alpha(u), u') \neq \overline{\Lambda}(\mathbf{P}_{\alpha+1}(u), u')$.*

Proof. The proof is similar to the proof of Lemma 40. Without mentioning it, we will repeatedly use Lemma 46 and the hypothesis of Theorem 45 that the cost functions are prefix-independent.

The first statement is immediate from definition of $\overline{\Lambda}$. For the second statement, consider $\rho = \rho_0\rho_1 \ldots \in \mathbf{E}_\alpha(v)$. Then by definition of $\mathbf{E}_\alpha(v)$, there exist $n \in \mathbb{N}$, $i \in \Pi$, and $u, u', v' \in V$ with $u' \neq v'$, $(u, v') \in E$, such that $\rho_n = u \in V_i$, $\rho_{n+1} = u'$, and

$$\forall \rho' \in \mathbf{P}_\alpha(v') : \lambda_i(\rho) > \lambda_i(\rho').$$

(See Figure 9 adapted to the context of Lemma 48). Let us prove that $\overline{\Lambda}(\mathbf{P}_\alpha(u), u') \neq \overline{\Lambda}(\mathbf{P}_{\alpha+1}(u), u')$. As $\rho \in \mathbf{P}_\alpha(v)$, then $\rho_{\geq n} \in \mathbf{P}_\alpha(u)$ by Lemma 22, which implies that $\Lambda_i(\mathbf{P}_\alpha(u), u') \geq \lambda_i(\rho_{\geq n})$. Let $\varrho \in \mathbf{P}_\alpha(u)$ be such that ϱ starts with edge (u, u') and has maximal cost $\Lambda_i(\mathbf{P}_\alpha(u), u')$. One gets

$$\lambda_i(\varrho) = \Lambda_i(\mathbf{P}_\alpha(u), u') \geq \lambda_i(\rho_{\geq n}) = \lambda_i(\rho) > \lambda_i(\rho').$$

Hence, using the same set $\mathbf{P}_\alpha(v')$ as for ρ, it follows that $\varrho \in \mathbf{E}_\alpha(u)$ for all such plays ϱ. Therefore $\mathbf{P}_{\alpha+1}(u) \subsetneq \mathbf{P}_\alpha(u)$ and $\Lambda_i(\mathbf{P}_\alpha(u), u') < \Lambda_i(\mathbf{P}_{\alpha+1}(u), u')$. ◀

As a consequence, the ordinal α_* of Proposition 17 is an integer. The proof is the same as for Corollary 41.

▶ **Corollary 49.** *There exists an integer α_* such that $\mathbf{P}_{\alpha_*}(v) = \mathbf{P}_{\alpha_*+1}(v)$ for all $v \in V$.*

As done for quantitative reachability games, let us now prove that the sets $\mathbf{P}_\alpha(v)$ are ω-regular for all α and v.

▶ **Lemma 50.** *Each $\mathbf{P}_\alpha(v)$ is an ω-regular set.*

Proof. The proof is similar to the proof of Lemma 43 and Proposition 42 (It is even simpler). Recall that as soon as $\mathbf{P}_\alpha(v)$ is empty, then $\mathbf{P}_\beta(v) = \emptyset$ for all $\beta \geq \alpha$.

Like in Lemma 43, we first prove that if $\mathbf{P}_\alpha(v)$ is MSO-definable, then $\Lambda_i(\mathbf{P}_\alpha(v))$ is computable for each $i \in \Pi$. Let ϕ be an MSO-sentence defining $\mathbf{P}_\alpha(v)$. One can decide whether $\mathbf{P}_\alpha(v)$ is empty. If this is the case, then $\Lambda_i(\mathbf{P}_\alpha(v)) = -\infty$ for all i. Suppose that $\mathbf{P}_\alpha(v) \neq \emptyset$, and let $i \in \Pi$ and $c \in C_i$. By hypothesis, the set of plays ρ in $(\mathcal{G}, v)$ with cost $\lambda_i(\rho) = c$ is ω-regular and thus MSO-definable by a sentence $\varphi_{c,i}$. We can thus decide whether $\mathbf{P}_\alpha(v)$ has a play ρ with cost $\lambda_i(\rho) = c$, thanks to sentence $\phi \wedge \varphi_{c,i}$. Therefore, by considering decreasing constants $c \in C_i$, we can decide whether $\Lambda_i(\mathbf{P}_\alpha(v)) = c$. This shows that $\Lambda_i(\mathbf{P}_\alpha(v))$ is computable.

[15] More precisely each component in $(\mathbb{Q} \cup \{-\infty\})^{|\Pi|}$ is restricted to $C_i \cup \{-\infty\}$.

Let us now prove that each set $\mathbf{P}_\alpha(v)$ is MSO-definable by induction on α. For $\alpha = 0$, we use the same defining MSO-sentence as in the proof of Proposition 42:

$$Q_v(0) \wedge \forall x \cdot \vee_{(u,u')\in E}(Q_u(x) \wedge Q_{u'}(x+1)).$$

Let $\alpha \in \mathbb{N}$ be a fixed integer. By induction hypothesis, each set $\mathbf{P}_\alpha(v)$ is MSO-definable, and $\Lambda_i(\mathbf{P}^I_\alpha(v))$, $i \in \Pi$, is computable by the first part of the proof. These sets and constants can be considered as fixed. The only case to consider is $\mathbf{P}_\alpha(v) \neq \emptyset$ (recall that this property is decidable). To show that $\mathbf{P}_{\alpha+1}(v)$ is also MSO-definable, it is enough to prove that $\mathbf{E}_\alpha(v)$ is MSO-definable. Recall that $\rho \in \mathbf{E}_\alpha(v)$ iff there exist $n \in \mathbb{N}$, $i \in \Pi$, and $u, u', v' \in V$ with $u' \neq v'$, $(u, v') \in E$, such that $\rho_n = u \in V_i$, $\rho_{n+1} = u'$, and $\forall \rho' \in \mathbf{P}_\alpha(v')$: $\lambda_i(\rho) > \lambda_i(\rho')$. The last condition can be replaced by $\lambda_i(\rho) > \Lambda_i(\mathbf{P}_\alpha(v')) \neq -\infty$.[16] Let us provide an MSO-sentence ψ defining $\mathbf{E}_\alpha(v)$:

$$\exists n \cdot \bigvee_{\substack{i\in\Pi, u\in V_i \\ u'\neq v'\in V \\ (u,v')\in E}} \quad \bigvee_{\substack{c\in C_i \\ c>\Lambda_i(\mathbf{P}_\alpha(v'))\neq-\infty}} (Q_u(n) \wedge Q_{u'}(n+1) \wedge \varphi_{c,i})\,.$$

◀

We get the next corollary. The proof is the same as for Corollary 44.

▶ **Corollary 51.** *If $\mathbf{P}_\alpha(v) \neq \emptyset$, then one can compute a lasso play $h \cdot g^\omega$ in $\mathbf{P}_\alpha(v)$ with $\lambda_i(h \cdot g^\omega) = \Lambda_i(\mathbf{P}_\alpha(v))$. This play depends on i and v.*

We are now able to prove the main result of this section.

Proof of Theorem 45. Let $(\mathcal{G}, v_0)$ be a game satisfying the hypotheses of Theorem 45. Let us summarize the results of the previous lemmas. We know that $\alpha_* \in \mathbb{N}$ and that one can construct the sets $\mathbf{P}_{\alpha_*}(v)$, $v \in V$. As these sets are ω-regular, one can decide whether they are all non-empty. In case of positive answer, there exists a weak SPE in $(\mathcal{G}, v_0)$ by Theorem 18. If in addition some constants $c_i \in C_i$ are given, then the set $\mathbf{P}_{\alpha_*}(v_0) \cap \{\rho \text{ in } (\mathcal{G}, v_0) \mid \lambda_i(\rho) \leq c_i, \forall i \in \Pi\}$ is also ω-regular. Hence one can also decide whether this set is non-empty and thus whether there exists a weak SPE in $(\mathcal{G}, v_0)$ with cost component-wise bounded by $\bar{c}$. This establishes the first part of Theorem 45.

Suppose that such a weak SPE exists, then let us show that we can construct a weak SPE that is finite-memory with the same construction as in the proof of Theorem 32. By Corollary 51, for all $i \in \Pi$, $v \in V$, one can construct a lasso play $h_{i,v} \cdot (g_{i,v})^\omega \in \mathbf{P}_{\alpha_*}(v)$ with maximal cost $\lambda_i(h_{i,v} \cdot (g_{i,v})^\omega) = \Lambda_i(\mathbf{P}_{\alpha_*}(v))$. The construction of a finite-memory SPE $\bar{\sigma}$ from the finite set of lasso plays $h_{i,v} \cdot (g_{i,v})^\omega$ is conducted as in the proof of Lemma 21. It is done step by step thanks to a labeling γ of the non-empty histories.

Initially, none of the histories is labeled. We start with history v_0 and with any play $h_{i,v_0} \cdot (g_{i,v_0})^\omega \in \mathbf{P}_{\alpha_*}(v_0)$, $i \in \Pi$.[17] The strategy profile $\bar{\sigma}$ is partially defined such that $\langle \bar{\sigma} \rangle_{v_0} = h_{i,v_0} \cdot (g_{i,v_0})^\omega$, and the non-empty prefixes h of $h_{i,v_0} \cdot (g_{i,v_0})^\omega$ are all labeled with $\gamma(h) = (i, v_0)$.

At the following steps, we consider a history $h'v'$ that is not yet labeled, but such that h' has already been labeled. By induction, $\gamma(h') = (j, v)$ and there exists $hv \leq h'$ such that

[16] Set $\mathbf{P}_\alpha(v')$ must be non-empty.
[17] When some constants $c_i \in C_i$ are additionally given, play $h_{i,v_0} \cdot (g_{i,v_0})^\omega$ must be replaced by any lasso play in $\mathbf{P}_{\alpha_*}(v_0) \cap \{\rho \text{ in } (\mathcal{G}, v_0) \mid \lambda_i(\rho) \leq c_i, \forall i \in \Pi\}$.

$\langle \bar{\sigma}_{\restriction h} \rangle_v = h_{j,v} \cdot (g_{j,v})^\omega$. Suppose that $\mathrm{Last}(h') \in V_i$, the proof of Lemma 21 requires to choose a play $\rho' \in \mathbf{P}_{\alpha_*}(v')$ such that $\lambda_i(h'\rho') = \lambda_i(\rho') \geq \lambda_i(h \cdot \langle \bar{\sigma}_{\restriction h} \rangle_v)$. We choose $\rho' = h_{i,v'} \cdot (g_{i,v'})^\omega$ with maximal cost $\lambda_i(\rho') = \Lambda_i(\mathbf{P}_{\alpha_*}(v'))$. Then we continue the construction of $\bar{\sigma}$ such that $\langle \bar{\sigma}_{\restriction h'} \rangle_{v'} = h_{i,v'} \cdot (g_{i,v'})^\omega$, and for all non-empty prefixes g of $h_{i,v'} \cdot (g_{i,v'})^\omega$, we define $\gamma(h'g) = (i, v')$.

We know by Lemma 21 that $\bar{\sigma}$ is an SPE. It is finite-memory because it only depends on the finite number of lasso plays $h_{j,v} \cdot (g_{j,v})^\omega$, and the labeling γ that can be computed inductively as follows. Initially, $\gamma(v_0) = (i, v_0)$ for some chosen $i \in \Pi$. Let $h' \in \mathrm{Hist}_i$ and suppose that $\gamma(h') = (j, v)$. If $h'v'$ respects $h_{j,v} \cdot (g_{j,v})^\omega$, then $\gamma(h'v') = (j, v)$, otherwise $\gamma(h'v') = (i, v')$ (as explained in the previous paragraph). This establishes the second part of Theorem 45. ◀

6 Conclusion and Future Work

In this article, we have studied the existence of (weak) SPEs in quantitative games. We have proposed a Folk Theorem for weak SPEs, and a weaker version for SPEs. To illustrate the potential of this theorem, we have given two applications. The first one is concerned with quantitative reachability games for which we have provided an algorithm to compute a finite-memory SPE, and a second algorithm for deciding the constrained existence of a (finite-memory) SPE. The second application is concerned with another large class of games for which we have proved that the (constrained) existence of a (finite-memory) weak SPE is decidable.

Future possible directions of research are the following ones. We would like to study the complexities of the problems studied for the two classes of games. We also want to investigate the application of our Folk Theorem to other classes of games. The example of Figure 3 is a game with a weak SPE but no SPE (see Example 12). Recall that for this game, the cost $\lambda_i(\rho)$ can be seen as either the mean-payoff, or the liminf, or the limsup, of the weights of ρ. We do not know if games with this kind of payoff functions always have a weak SPE or not.

References

1. Thomas Brihaye, Véronique Bruyère, Julie De Pril, and Hugo Gimbert. On subgame perfection in quantitative reachability games. *Logical Methods in Computer Science*, 9, 2012.
2. Thomas Brihaye, Julie De Pril, and Sven Schewe. Multiplayer cost games with simple Nash equilibria. In *LFCS*, volume 7734 of *Lecture Notes in Comput. Sci.*, pages 59–73. Springer, 2013.
3. Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Secure equilibria in weighted games. In *CSL-LICS*, page 26. ACM, 2014.
4. Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. *ACM Trans. Comput. Log.*, 11, 2010.
5. János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan, and Koos Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Math. Oper. Res.*, 35:742–755, 2010.
6. Drew Fudenberg and David Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31:251–268, 1983.
7. Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.
8. Miroslav Klimos, Kim G. Larsen, Filip Stefanak, and Jeppe Thaarup. Nash equilibria in concurrent priced games. In *LATA*, volume 7183 of *Lecture Notes in Comput. Sci.*, pages 363–376. Springer, 2012.

9 H.W. Kuhn. Extensive games and the problem of information. *Classics in Game Theory*, pages 46–68, 1953.

10 Jean-François Mertens. Repeated games. In *Internat. Congress Mathematicians*, American Mathematical Society, pages 1528–1577. Providence, RI, 1987.

11 J. F. Nash. Equilibrium points in n-person games. In *PNAS*, volume 36, pages 48–49. National Academy of Sciences, 1950.

12 Martin J. Osborne and Ariel Rubinstein. *A course in Game Theory.* MIT Press, Cambridge, MA, 1994.

13 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.

14 Stéphane Le Roux and Arno Pauly. Infinite sequential games with real-valued payoffs. In *CSL-LICS*, page 62. ACM, 2014.

15 Ariel Rubinstein. Comments on the interpretation of game theory. *Econometrica*, 59:909–924, 1991.

16 Eilon Solan and Nicolas Vieille. Deterministic multi-player dynkin games. *Journal of Mathematical Economics*, 39:911–929, 2003.

17 Wolfgang Thomas. Automata on infinite objects. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B)*, pages 133–192. 1990.

18 Michael Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *FSTTCS*, volume 4337 of *Lecture Notes in Comput. Sci.*, pages 212–223. Springer, 2006.